



UNIVERSITÀ
DEGLI STUDI
DI UDINE

Università degli studi di Udine

Personalized Egocentric Video Summarization of Cultural Tour on User Preferences Input

Original

Availability:

This version is available <http://hdl.handle.net/11390/1124980> since 2022-05-16T12:26:50Z

Publisher:

Published

DOI:10.1109/TMM.2017.2705915

Terms of use:

The institutional repository of the University of Udine (<http://air.uniud.it>) is provided by ARIC services. The aim is to enable open access to all the world.

Publisher copyright

(Article begins on next page)

Personalized Egocentric Video Summarization of Cultural Tour on User Preferences Input

Patrizia Varini, Giuseppe Serra, and Rita Cucchiara, *Member, IEEE*

Abstract—In this paper we propose a new method for customized summarization of egocentric videos according to specific user preferences, so that different users can extract different summaries from the same stream. Our approach, tailored on cultural heritage scenario, relies on creating a short synopsis of the original video focused on key shots, in which concepts relevant to user preferences can be visually detected and the chronological flow of the original video is preserved. Moreover, we release a new dataset, composed of egocentric streams taken in uncontrolled scenarios, capturing tourists cultural visits in six art cities, with geolocalization information. Our experimental results show that the proposed approach is able to leverage user's preferences with an accent on storyline chronological flow and on visual smoothness, since the produced summaries have been preferred by 85% of users.

I. INTRODUCTION

In recent years the widespread use of wearable cameras allows capturing everyday life activities such as sport, education, social interactions and cultural heritage visits. However, these videos typically consist of long streams of data with a ceaseless unstable appearance and frequent changes of observer's focus. Moreover, differently from the edited videos, they lack of cuts to separate different video segments. Therefore, there is a urgent demand of adequate tools to process the semantic video content in order to select and store only the meaningful sequences.

Based on this demand, a large number of different egocentric video summarization techniques has been proposed recently. Several approaches have explored low-level, mid-level and high-level features [1]–[3]. Low-level features, typically related to visual motion and color, are able to identify the most stable and homogeneous sequences; the mid and high-level ones [4]–[6] are useful to identify salient content such as social interactions or objects manipulations. Furthermore, multimedia approaches have also investigated physiological data such as brain waves [7] and gaze [8] to select the relevant segments.

Although these summarization techniques deal with egocentric characteristics they produce a unique summary. The generation of a single summary could be adequate in some controlled domains such as video surveillance of a specific area [9] [10], in which the salient events are generally pre-defined. This assumption can not be exploited in egocentric videos captured in unconstrained scenarios. In this case users

can have different preferences. Therefore, they might prefer retaining in the summary some events rather than others. For instance, art lovers and fashion enthusiasts might desire to extract different summaries from their egocentric videos taken during a visit in an art city. In addition, a user may want to share personalized summaries with different people, such as parents or friends.

In this paper we address this challenge. In particular, we propose a summarization technique in the cultural heritage domain that is able to generate customized summaries according to specific user preferences. To identify candidate relevant sequence in an egocentric video, we first present a behavior pattern classifier based on a novel CNN architecture exploiting apparent, 3D motion and visual assessment features. Based on these detected sequences, the proposed summarization approach selects the items considering three key elements: the attention behavior of the wearer, the visual and semantic coherence with the user preferences and a narrativity grade gauged by a new proposed metric.

To assess sequences related to the user preference, we propose to build a set of visual classifiers. Since preferences are potentially limitless, semantic classifiers are built on-fly using a data-driven approach, that exploits geolocalization and DBpedia semantic knowledge. Moreover, to recover the chronological flow of the most relevant sequences we propose a Personalized Page Rank approach that takes into account narrativity metric assessed by means of importance and relevance into the frame sequence. To the best of our knowledge, this is the first attempt in egocentric vision summarization to deal with the user intentions in order to generate personalized summaries.

Finally, we present several experiments on a new dataset of unconstrained egocentric videos and we demonstrate the effectiveness of our solution. Figure 1 shows an example of three different summaries: an uniform summary no preference driven, and two customized summaries obtained with the proposed method, using two sets of different user preferences.

To sum up, our contribution can be summarized as follows:

- we propose an user behavior classifier, based 3D CNN approach, that joints apparent, 3D real motion features, that achieves state of the art results;
- to identify the sequences related to user preferences we present a novel data-driven approach based on a set of classifiers dynamically learned exploiting DBpedia and geolocalization of the user;
- we define a new technique to recover in the summary the original chronological storyline progression of the original video;



Fig. 1: Different summaries produced by the same original video, one uniform summary and two customized summaries based on our method according to two different sets of user preferences, “Sport Cars”, “Romanesque” (related frames edges surrounded respectively in blue and green) and “Shopping, Fashion” (related frames edges surrounded in red).

- we publicly release a new unconstrained dataset containing egocentric videos of cultural heritage visits in six art cities. To best of our knowledge, this is the first egocentric video dataset that contains jointly geolocalization information, behavior classification, narrative milestones.

II. RELATED WORKS

Although video summarization has been studied in different application domains such as sport [11]–[13], surveillance [14], [15], multimedia retrieval [16], [17], only recently first-person summarization has been addressed by research community. Differently from the third-person approaches, egocentric vision summarization requires new techniques to deal with continuous changes of observers focus and with lack of hard cuts between scenes.

To address these challenges, several egocentric approaches have explored the usage of low-level and high-level features such as color, visual motion and presence of people or objects [18], [19].

In particular, Lu et al. [1] presented a technique based on optical flow and blurriness features to divide videos into sub-shots. Subsequently, they selected a chain of sub-shots choosing the ones where reciprocal influence between important objects and characters was detected.

Gygli et al. [5] identified video segments by using a “superframe” segmentation on different raw videos, then they scored visual interestingness per superframe using a set of low-features (color, edges, blur), mid-features (landmarks) and high-level concepts (recognition of faces, objects using segmentation). Based on these scores, they selected an optimal subset of superframes to create a maximum informative synopsis. Moreover, Gygli et al. [20] formulated the task of video summarization as a subset selection problem. In particular, they learned a linear combination of submodular functions, using structured learning with a large-margin formulation, directly from reference summaries created by human annotators.

Lidon et al. [21] first selected informative images by a two-class CNN. Then, relevance ranking was obtained by integrating techniques for saliency detection, object recognition and face detection. Poleg et al. [3] learnt to classify

camera’s wearer motion leveraging a three dimensional CNN architecture, fed with stream sparse optical flow volumes as input. In addition, Halperin, Poleg et al. [2] proposed a fast-forwarding approach based on motion related features that obtained timelapses or hyperlapses by selecting optimal set of frames avoiding non-aligned consecutive sampled frames.

Joshi et al. [22] obtain hyperlapses by choosing frames from the original video that best match a certain speed-up, surreptitiously achieving camera motion smoothing. Okamoto et al. [23] starting from first person videos produces short walking route guides in which they dynamically control video playing speed. This is achieved assessing egocentric motion features assigning different importance score to the sequence of the streams in which appear urban items such as pedestrian crosswalks or traffic lights. Lin et al. [24] propose a Structured SVM detector to detect highlights analyzing the context of each video sequence. Yao et al. [25] learned temporal and spatial highlights from appearance of video frames and temporal dynamics across frames using late fusion of two-streams spatio-temporal CNN structure. Zhang et al. [26] investigated how to apply Long Short-Term Memory model to supervised video summarization.

To improve the summaries accuracy, other approaches presented solutions based on information extracted from heterogeneous sensors. For instance, Ng et al. [27] proposed an egocentric video summarization method that used physiological data such as brain waves to select the relevant segments of video. Lee et al. [4] proposed an approach that focused on the detection of important cues for each frame, such as objects and people the camera wearer interacts with. To identify these frames they proposed to combine wearers gaze and visual features, such as objects’ motion and appearance. More recently, Xu et al. [8] used gaze information to identify the skims and to score their relevance, subsequently maximized reciprocal shots information in order to achieve a summary in which relevance, diversity, fidelity and compactness were optimized.

Although these egocentric summarization approaches are valuable, they miss a key element of the egocentric paradigm: the person’s specific intention. In fact, everyone is different

thus there doesn't exist a univocal summary which is the best one for everybody.

To address this issue, an early version of the proposed approach was introduced in a preliminary work [28]. Here, we propose new key components in order to achieve more reliable results. Differently from our previous work, we present a new algorithm for Behavior Pattern Recognition that is able to achieve robust results even in unconstrained scenarios. We propose a more formal approach, based on DBpedia, for the semantic expansion and filtering of the preferences specified by the user. We introduce a new solution to preserve the chronological flow in the summary. Finally, we create, and contextually publicly release, a larger egocentric video dataset of Cultural Heritage tour visits.

III. PERSONALIZED EGOCENTRIC VIDEO SUMMARIZATION

Our solution takes a long first-person video of cultural or touristic experience and a text expressing user preferences as input and returns a customized short video summary as output. Our goal is to preserve logical and chronological storyline in which we want to put in evidence the segments related to the preferences expressed by a particular user. Figure 2 synthesizes our approach.

In egocentric videos of cultural heritage experiences, the presence of relevant information is likely to be put in relation with camera's motion behavior patterns more than others. For example, an interesting behavior "looking at an artistic monument" is characterized by some little movements of the head for an interval of time sufficient to establish a visual fixation pattern [29], probably standing still for a while. This behavior has a very different pattern of motion with respect to other less interesting behaviors, like wandering around or walking straight that show wide movements of the head.

Therefore, we define a set of motion behavior classifiers based on real, apparent motion and quality assessment of the visual content to extract homogeneous skims of the video (from now on we call them "ego-shots" - see Section III-A for more details). Using behavior detection we are also able to assign a score to ego-shots according to the observer's behavior pattern: we assign a high score to ego-shots with the presence of user attention and lower scores to ones containing less relevant activities such as wandering, walking, running, etc.

However, recognizing the presence of attention behavior pattern alone is not sufficient. In fact, actions like "looking at the phone", "waiting to cross a road", "looking at bus timetable", may all be put in relation with attention motion behavior pattern, but they are generally uninteresting in this scenario. Consequently, we introduce in our solution two other key aspects: "semantic relatedness" with the user preferences and "narrative importance" to recover original storyline.

Semantic relatedness is achieved building a set of visual classifiers based on the user preferences. However, since topics of interest are potentially unlimited, classification based on a number of predefined rigidly defined classes may lead to poor performances. Thus, we build specific classifiers for the

topics of interest using a data-driven approach that extracts on-fly reliable training images exploiting DBpedia semantic knowledge and GPS tour information.

Narrative importance has the target of recovering the chronological flow of most salient events. For this purpose we want to identify the pivotal skims for each sequence in the narration, as the segments with the highest amount of information.

Let V be the original video regarded as the composition of a set of consecutive shots $\{S_1, S_2, \dots, S_n\}$, the personalized video summarization is arranged in chronological order, selecting the ego-shots S^* that maximize the objective function depending from Behavior score $B(S)$, User Preference Semantic Score $P(S)$ and Narrative Importance score $N(S)$:

$$O(S^*) = \text{Arg Max}_{S_i \in V} (B(S_i) + P(S_i) + N(S_i)) \quad (1)$$

It should be noted that we could adopt a weighted sum between the three scores in order to gauge the contributions. However, while a optimization of these weights could slightly improve the performance, it would lead to a fine tuning too excessively specific in contrast with our purposes.

A. Behavior Pattern Detection

Tailored on a typical cultural experience scenario, we assume the hypothesis that shaped behavior of visitors motion is a relevant information to identify and classify interesting or irrelevance sequence of the video. Thus, we first segment a video into ego-shots characterized by homogeneous behavior patterns. In detail, we define six behavior pattern classes: "Attention" (where user is paying attention to something), "Looking around" (searching for something etc), "Walking", "Running", "On wheels", "Wandering".

To detect these activities we propose a CNN with 10 layers trained on motion and frame visual assessment features (see Fig. 3). Motion features encode both visual apparent and real components. Visual motion feature is based on optical flow estimated using the Farneback algorithm [30]. In particular, considering the optic flow computed for each couple of consecutive frames, we extract the relative apparent velocity and acceleration gradient for both horizontal and vertical components for each pixel.

Real motion features are obtained collecting the geographic locations of the visitor by means of a GPS sensor. Locations are represented as a temporal sequence of the spatio-temporal points, meant as pairs compound with coordinate in space and in time $\{p_0 = (x_0, y_0, z_0, t_0), \dots, p_N = (x_N, y_N, z_N, t_N)\}$, where $(x_i, y_i, z_i) \in \mathcal{R}^3, t_i \in \mathcal{R}^+$ for $i = 0, 1, \dots, N$ and $t_0 < t_i < t_N$. Given them, we extract three components of real 3D velocity of visitor as:

$$V_x = \frac{|x_{i+1} - x_i|}{t_{i+1} - t_i}, V_y = \frac{|y_{i+1} - y_i|}{t_{i+1} - t_i}, V_z = \frac{|z_{i+1} - z_i|}{t_{i+1} - t_i}. \quad (2)$$

To assess the frame quality, we compute frames blur feature by using the method in [31]. This approach assumes that the sharpness of a frame is contained in its gray component and estimate the blur annoyance only on the luminance component, computing and evaluating the line and row difference between

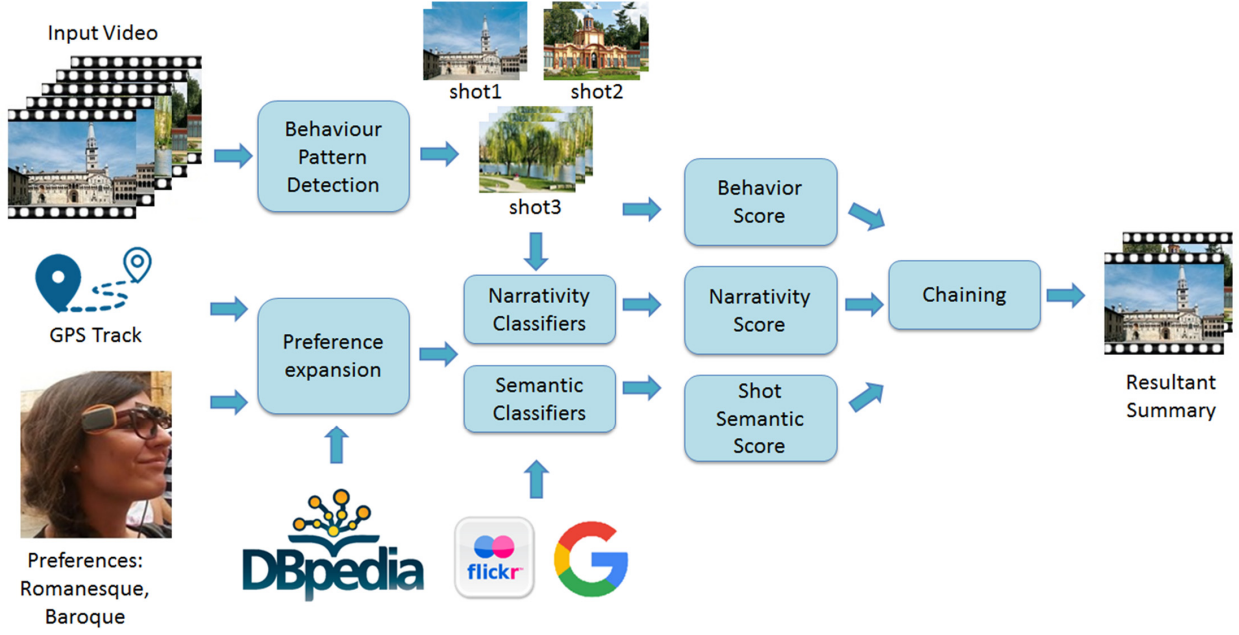


Fig. 2: Schematization of the proposed method: maximization of objective function based on visitor behavior, semantic and narrativity scores is computed on detected shots to obtain final summary.

the original image and the image obtained applying to it a horizontal and a vertical strong low-pass filter. The blur feature employed in our model is the difference between the original image and the filtered image.

Apparent motion and blur descriptors (originally shaped as the frame size) are resized to 96×96 pixel and stacked to obtain our overall apparent motion descriptor for a couple of frames, a tensor shaped $96 \times 96 \times 5$. To take into account time dimension the tensor descriptors are further stacked over one second that provides a final descriptor of shape $96 \times 96 \times 5 \times 25$ (resampling to 25 fps videos whose frame rate is different).

Our network can be described in shorthand notation as C1(64; 11; 3; 3)-P1-C2(128; 5; 2; 1)-P2-C3(256; 3; 2; 1)-C4(384; 3; 2; 1)-C5(256; 3; 2; 1)-P5-FC6(2048)-FC7(1024), with 3 dropout layer, 0.1, 0.3, 0.5, during training phase. In the above notation $C(d, f, s, p)$ indicates a convolutional layer with d filters, kernel size $f \times f$, depth p , applied to the input with stride s . The first convolutional layer and pooling layer are implemented as 3D, while following ones are standard 2D. FC(z) is a fully connected layer with z nodes. Rectified Linear Units (ReLU) is employed. We adopt the L2-SVM loss to train the CNN [32].

The L2-SVM loss function is differentiated w.r.t. the activation \mathbf{h} of the penultimate layer FC7(z). This loss function $L(\mathbf{w})$ can be written in terms of the weights matrix \mathbf{w} and the input of the penultimate layer \mathbf{h} , as:

$$\frac{\partial L(\mathbf{w})}{\partial \mathbf{h}_n} = -2P\theta_n \mathbf{w}(\max(1 - \mathbf{w}^T \cdot \mathbf{h}_n, 0)), \quad (3)$$

where $P \in \mathbb{R}$ is the model penalty parameter and $\theta_n \in \{-1, +1\}$.

Once the aforementioned CNN is trained from scratch on the training set, FC7 activations are extracted and concatenated

with real motion features, after L2 normalization. A linear SVM is then used to classify shots of the egocentric sequences. Figure 3 presents a schematization of our behavior classifiers.

B. Semantic classifiers based on User Preferences

To score the semantic relatedness of the ego-shots to the User Preferences, we build a recognition system based on visual classifiers. Since topics of interest requested by the user can be potentially limitless, we propose a data-driven approach that gathers positive and negative training images from the web.

In our solution, for each video the user can indicate multiple preferences in the form of a set of sequences of words. Given a user preference a pre-processing step to eliminate stopwords, dates, numbers, punctuations and symbols, is performed; it is converted in the lemmatized form. Thus, a single user preference can be represented as a set of lemmatized concepts $C = \{c_1, c_2, \dots, c_n\}$.

To extract positive and negative training images from the web, it is necessary to define a set of adequate text input queries. In fact, a mere query using only the concatenation of user preference terms might achieve poor results due to language and semantic ambiguity.

Since Cultural Heritage items are often highly location-specific and since some concepts and semantic relationships may be unknown even to the user, we show that semantic exploration on a reliable Knowledge Base (KB) and user localization (L) can improve the search terms to collect training images.

To achieve this target, we choose to exploit DBpedia Knowledge Base, because it is an important and constantly updated dataset in Linked Open Data, that consists of semantic struc-

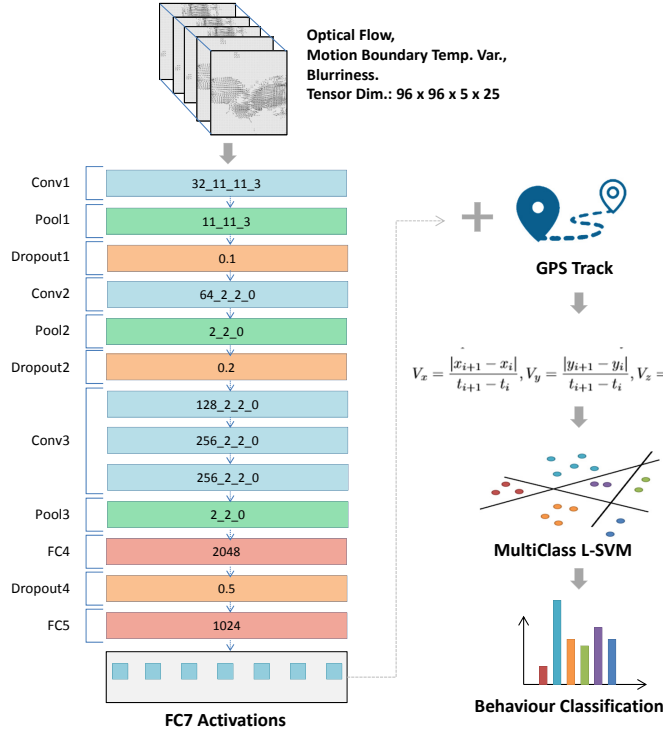


Fig. 3: Behaviour classification schema. Dropout layers are used only during training phase. Convolutional layer parameters are expressed as $d_f_s_p$, where d is the number of filters, $f \times f$ is the kernel size, p is their depth, s the stride applied to the input. Pooling layers parameters are expressed as f_s_p , where $f \times f$ is the kernel size, s is the stride, p the depth.

tured content (hyponyms, hyperonyms, synonyms, antonyms) with cross references between related topics [33].

Dbpedia can be represented as an undirected graph $KG = \{V, E\}$ where V are the nodes representing stemmed and lemmatized concepts and $E \subset V \times V$ are the edges representing the links among nodes.

For each of concept c_i of a user preference we attempt to match it with Dbpedia graph KG using a text match. In the case of exact match, there might be the need of disambiguation with nodes of KG . If disambiguation is needed, we choose the node that minimizes the sum of Dijkstra distances [34] between the ambiguous node and nodes related to other concepts of C and location L .

In this way, it is possible to associate concepts in C to KG nodes $K = \{k_1, k_2, \dots, k_m\}$, which we call “Seed nodes” (notice $m = n$ in the case in which all the concepts are matched to KG).

Therefore, a subgraph from KG considering only the set of nodes K and their edges is built. This subgraph is enriched recursively taking into account, for each k_i , all the nodes connected by edges related to Cultural Heritage domain (i.e. type skos:broader, skos:broaderOf, rdfs:sub-ClassOf, rdfs:type and dcterms:subject).

For what concerns the number of recursive steps, based on our experiment we found that a distance of two-hops is

optimal in most cases. In fact, node expansions to three hops usually produce very large graphs and introduce a considerable amount of noise (see Section IV for more details). Concepts in C , that are not linked to any nodes of KG , are not expanded. For example, let’s consider the simple user preference $C = \{\text{“Romanesque”, “Church”}\}$; the corresponding set of seed nodes on the KG is $K = \{\text{“Church”, “Romanesque Architecture”}\}$.

Once obtained the graph, to minimize the noise induced by concepts less related on the user preference, we detect and analyze the semantic communities to identify the most connected one centered on K nodes.

For this purpose we adopt the recursive Girvan-Newman algorithm [35]. This starts with computing the “betweenness” score for each of the edges (“betweenness” of an edge is the number of shortest paths between pairs of nodes that run along it). Then edges with the highest score are removed and the betweenness of all edges affected by the removal are computed. The last two steps are repeated until no edges remain.

We use graph metrics to assess how focused and how dense the detected communities around the expanded user preferences are. In particular, Focused Current Flow Betweenness Centrality [36] (FCFBC), Average Node Connectivity [37] (AVC), Focused Information Centrality [38] (FIC) and Average of the shortest paths between community nodes and location L (computed using Dijkstra algorithm), are adopted, in order to evaluate density of the communities and their semantic distance from the location.

At last starting from $K = \{k_1, k_2, \dots, k_m\}$ we identify a new set of concepts $\hat{K} = \{k_1, k_2, \dots, k_m, k'_1, k'_2, \dots, k'_p\}$, where $k'_1 \dots k'_p$ are the additional nodes of the detected community.

The Figure 4 shows the graph obtained by recursively expanding seed nodes of K and the detected communities. For sake of clarity we show only a bunch of nodes since the cloud of connected nodes is much larger.

Starting from \hat{K} , we download images and their metadata from Visual Online Repositories like Flickr, Google Photos, leveraging textual queries to extract positive and negative samples for each user preference.

The query to extract a set of positive image samples P , is obtained by combining in “OR” each element. For each image in P we extract a set of metadata, which consists of tags and/or description (preprocessed eliminating stopwords and applying stemming).

The query, to extract negative samples, is generated by combining in “NOR” the set of hyperonyms of \hat{K} nodes.

To further reduce noise and increase reliability of our training sample images we leverage jointly visual and semantic features extracted by textual metadata (tags and descriptions). To extract and encode visual content of each image, we use the convolutional network VGGNET-16, pre-trained on Imagenet, that has shown reliable results in image classification [39]. In particular, we use the VGGNET-16 as a feature extractor removing the output layer and using the activations of the last fully connected layer, after applying L_2 normalization.

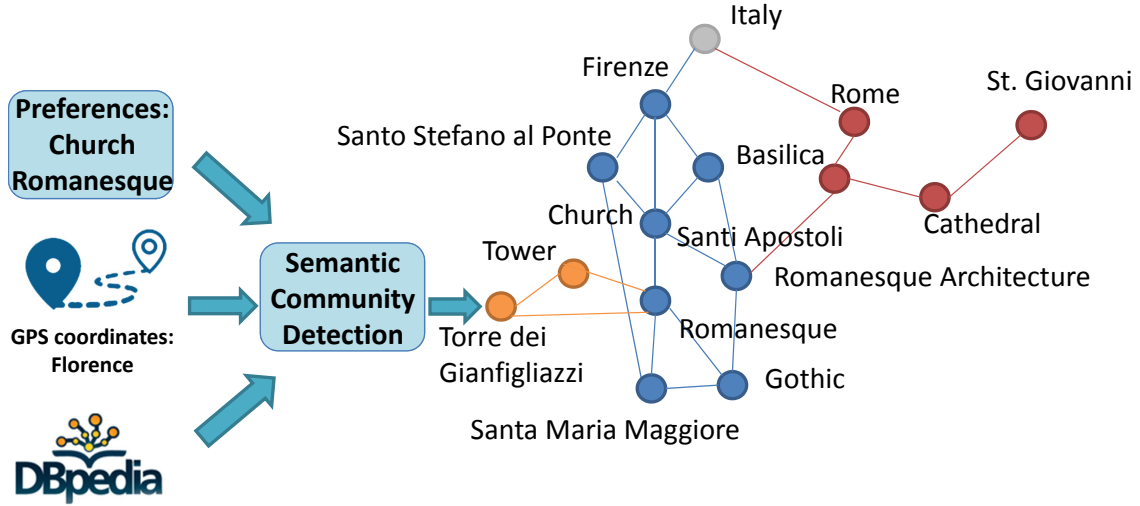


Fig. 4: Schematization of the semantic expansion approach: the Blue graph is the community with the maximum connectivity and centrality of the expanded graph; the Yellow and Red communities contain the nodes linked with the Seed nodes K but are less connected and less focused than the blue one.

Semantic representation features are extracted from downloaded image tags and descriptions using the Word2Vec algorithm [40]. We train Word2Vec model on DBpedia english abstracts, setting the dimensionality of our feature vector to 100. For each preprocessed metadata of a downloaded image I we compute its representation in the semantic vector space; then the final semantic feature representation \mathbf{w}_I of I is obtained by element-wise addition of their vector representations.

To filter the most noisy candidate training images, we discard the candidates that are semantically too far from the expanded concepts \hat{K} . To compute their semantic distance, we map \hat{K} in the semantic vector representation $\mathbf{w}_{\hat{K}}$ following the same approach as described above. Thus, semantic distance of I from original \hat{K} concepts, is obtained by computing the cosine distance between \mathbf{w}_I and $\mathbf{w}_{\hat{K}}$:

$$SS(\mathbf{w}_I, \mathbf{w}_{\hat{K}}) = \frac{\mathbf{w}_I^T \cdot \mathbf{w}_{\hat{K}}}{\|\mathbf{w}_I\| \|\mathbf{w}_{\hat{K}}\|} \quad (4)$$

Images with a semantic score $SS(\mathbf{w}_I, \mathbf{w}_{\hat{K}})$ below a threshold (empirically, ones with less than 10% of the average) are discarded.

For the non discarded images, we concatenate their metadata semantic representations to visual feature vector, obtaining a joint visual and semantic vectorial representation. Then, we perform unsupervised agglomerative clustering (DBSCAN, with dynamically adaptively computed parameters [41]) on the final features, discarding the images belonging to clusters with a number of elements less than 3% of the original image set. The non discarded images can now be used to build for each user preference a visual classifier. At last, for each shot we compute the User Preference Semantic Score $P(S)$ obtained averaging visual classification score over the shot length.

C. Shot narrativity score

At this point, according to Equation 1, we want to assign a score that quantifies narrative relevance of each ego-shot,

i.e. the capacity of a chosen subset of video segments to represent an important milestone in the chronological flow of events, with a focus on relevance and diversity. In fact, let's for instance assume that a visitor is walking for a while, far from any Points of Interest; POI detectors give to these scenes very low scores, but their complete elimination might generate a hole in the chronological flow. Recent approaches pursue mutual information maximization between original video and summary [8] assessing shots information diversity with methods like maximum margin relevance [42] or determinantal point process [26], eventually using people or objects as plot pivots [21]. Unlike egocentric videos of daily activities (a food preparation, object manipulation etc.) or videos taken by police body-worn cameras, typical Cultural Heritage tour egocentric videos consist of long sequences of events in hugely unconstrained scenarios. Thus, maximizing mutual information between summary and video or maximizing visual diversity of skims, lead to poor results, because a number of irrelevant but visually rare skims may obtain a passing score and might be included in the summary. Differently, we are interested in identifying the story most meaningful and influential joints in terms of "Relevance" and "Diversity".

To address this issue, we start regarding the original video V as a directed graph, where ego-shots $\{S_1, S_2, \dots, S_n\}$ are represented by nodes, edges and their weights are dependent on Relevance within a frame of time (see for more details subsection III-C1) and we then compute nodes centrality to assess their representativity. In fact, nodes centrality is widely used to select the most important nodes in terms of influence in a graph.

In particular, we propose to use the Personalized Page Rank (PPR) approach [43] to compute centrality. The Personalized Page Rank formulation yields a node centrality measure on weighted directed graphs, that depends on the number of received inbound links, the centrality of the linkers and, inversely, on their link tendency (the more lavish, the less

evaluated their contribution). In matrix form, the formulation for node centrality score can be written as:

$$\mathbf{C} = [d\mathbf{U} + (1 - d)\hat{\mathbf{W}}]^T \mathbf{C} \quad (5)$$

where matrix $\hat{\mathbf{W}}$ is obtained from the weighted graph adjacency matrix \mathbf{W} dividing by row sum: $\hat{W}(u, v) = \frac{W(u, v)}{\sum_{z \in \text{adj}(u)} W(u, z)}$, \mathbf{U} is a square matrix with all elements being equal to $1/n$, being n the total number of nodes in the graph (i.e. the number of ego-shots), and d is a “damping factor” fixed to 0.15 (based on preliminary results). Let’s be $\mathbf{Z} = [d\mathbf{U} + (1 - d)\hat{\mathbf{W}}]$ the Equation 5 can be written as:

$$\mathbf{C}^T = \mathbf{Z} \mathbf{C}^T. \quad (6)$$

This formulation shows that \mathbf{C}^T is the left eigenvector of the matrix \mathbf{Z} with the corresponding eigenvalue of 1. Based on Perron-Frobenius theorem [44] this eigenvector exists and can be uniquely computed. We solve this eigenvector equation using the Power Iteration method [45].

Now, sorting in descending PPR order the graph nodes, we re-score them penalizing their temporal proximity. At the end of the ranking process, the score values are normalized. The figure 5 synthetizes this approach.

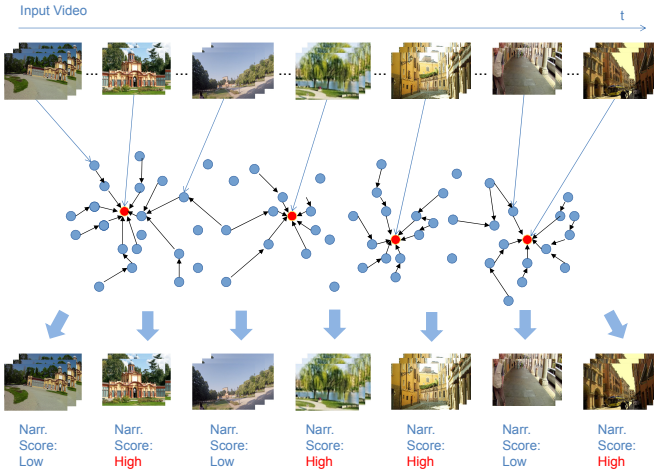


Fig. 5: Narrativity detection with Personalized Page Rank approach based on Visual Similarity and Visual Relevance. Narrativity score here is denoted as Low or High, instead of a numeric value, to make the example more readable.

1) *Weighted Adjacency Matrix definition*: we define edge weights between S_i and S_j as:

$$w_{i,j} = e^{-\frac{\Delta(N_{i,j})^2}{\tau}} \cdot P(S_i, S_j). \quad (7)$$

The term $e^{-\frac{\Delta(N_{i,j})^2}{\tau}}$ takes into account temporal distance between shots. In particular, $\Delta N_{i,j}$ is the difference in seconds between the shots S_i and S_j , and τ is a normalization factor (fixed empirically based on preliminary experiments). The factor $P(S_i, S_j)$ encodes visual similarity and relevance between two shots. We define it as:

$$P(S_i, S_j) = VS(S_i, S_j) \cdot \max(VR(S_i), VR(S_j)) \quad (8)$$

where $VS(S_i, S_j)$ is the function that measures the visual similarity between shot S_i and S_j and $VR(S_i)$ measures the visual relevance of the shot S_i . Edge direction is given by $\frac{R(S_i) - R(S_j)}{|R(S_i) - R(S_j)|}$ except in the case in which $R(S_i) = R(S_j)$; in this case the nodes are linked by two symmetric edges.

We compute visual similarity between ego-shots extracting features from deep 3-dimensional convolutional network named C3D [46] and computing cosine similarity between the last dense layer activations.

To compute visual relevance (VR) we assess three specific visual attributes: Colorfulness, Informativeness, Aesthetic Appearance. Indeed in Cultural Heritage domain, visitors are likely to spend long times engaged in potentially irrelevant activities involving colorless, dull backgrounds (staring at phones, informative boards, timetables), repetitive patterns (asphalt roads, brick walls, pavements, facades), crooked shots (odd shots captured during actions involving wide head movements).

To assess jointly the three attributes we exploit a multi-column CNN [47], that combines various CNN columns.

The single column CNN consists of the aforementioned C3D. Features are extracted from the last dense layer. To exploit the multiple attributes approach, multi-task learning is used [48], where feature representation and classification error minimization are constructed in a joint manner for the three attributes. Assuming we have Informativeness labels \mathbf{y}_I , Colorfulness labels \mathbf{y}_C and Aesthetic Appearance \mathbf{y}_A for all training images, the problem could be formulated as minimization of the cross-entropy loss:

$$\begin{aligned} L(\mathbf{x}, \mathbf{y}, \mathbf{w}) = & - \left(\sum_{i=1}^N \sum_{l \in L_C} \mathbf{I}(y_{Ci} = l) \ln p(y_{Ci} = l | x_i, \mathbf{w}_C) + \right. \\ & \sum_{i=1}^N \sum_{l \in L_I} \mathbf{I}(y_{Ii} = l) \ln p(y_{Ii} = l | x_i, \mathbf{w}_I) + \\ & \left. \sum_{i=1}^N \sum_{l \in L_A} \mathbf{I}(y_{Ai} = l) \ln p(y_{Ai} = l | x_i, \mathbf{w}_A) \right) \end{aligned} \quad (9)$$

where \mathbf{x} is the feature vector of N training images, L_C , L_I and L_A are the label sets. The sets $\mathbf{w}_I = \{\mathbf{w}_{I,l}\}_{l \in L_I}$, $\mathbf{w}_C = \{\mathbf{w}_{C,l}\}_{l \in L_C}$ and $\mathbf{w}_A = \{\mathbf{w}_{A,l}\}_{l \in L_A}$ are the model weights respectively for Colorfulness, Informativeness and Aesthetic Appearance, initialized to the values of the pre-trained single networks. $\mathbb{I}(z) = 1$ if z is true and 0 if false and, for each attribute a in {Colorfulness, Informativeness, Aesthetic Appearance}, probability p_a is expressed as:

$$p_a(y_{ai} = l | \mathbf{x}_i, \mathbf{w}_{al}) = \frac{e^{\mathbf{w}_{al}^T \mathbf{x}_i}}{\sum_{h \in L_a} e^{\mathbf{w}_{ah}^T \mathbf{x}_i}} \quad (10)$$

At last, for each shot, Relevance $VR(S_i)$ is computed averaging the prediction on the shot frames.

IV. EXPERIMENTAL RESULTS

To evaluate the performance of our approach we collected 48 videos captured by tourists during visits in six italian art

cities: Bologna, Ferrara, Modena, Ravenna, Reggio-Emilia, Parma. They show visitors experiences such as a visit to a cultural heritage site or an attraction (churches, monuments etc), or activities like shopping or walking. The cameras are placed on the tourist's head. Videos are of variable length, ranging from a minimum of 15 minutes to a maximum of 25 minutes taken in a uncontrolled setting with a resolution variable from 720×576 to 1920×1080 . Granularity of GPS sensor is one second in time and 1 meters linear displacement in space. To the best of our knowledge this is a novel dataset specifically designed to analyze this egocentric scenario ¹.

A. Experiments on Behavior Pattern classification

First, we examine the effectiveness of our solution for behavior recognition based on frame quality assessment, visual and real motion pattern. A subset of 26100 annotated frames of our dataset is fully annotated with the six behavior pattern classes and used to test our methodology. We use a 10-fold cross-validation approach partitioning the dataset for training, validation and testing, (dimensioned 80, 10 and 10 percentage of the dataset respectively). We compare our technique with recent and related approaches [1], [3], [28]. To be comparable for all of these methods we use the same experimental settings. Table I presents comparative results between our solution and these approaches in terms of average class F1-score. As can be seen, our solution achieves a better performance. This is mainly due to the addition, differently from previous approaches, of motion boundary and 3D real motion features in our 3D CNN architecture.

In particular, Figures 6 shows the confusion matrices for the aforementioned approaches. Adding descriptors related to 3D real motion helps to distinguish higher speed motions; this can be seen in classes such as “Running” and “Moving on Wheels”. Furthermore motion boundary features better encode the wide head motion or sprawl, improving the performance of classes such as “Looking around” and “attention” patterns.

TABLE I: Comparison of classification accuracy.

Method	Accuracy
Lu et al. [1]	63.67
Varini et al. [28]	74.17
Poleg et al. [3]	90.67
Our method	92.17

B. Experiments on User Preferences Expansion

To test the performance of our semantic expansion of the user preferences we simulate and analyze different user inputs. Table III shows the initial user preferences, the city locations and the expanded preferences. For each expanded item we evaluate the precision (i.e. the number of correct terms with respect to the number of expanded terms). The table also reports the measured Average Node Connectivity, Focused Information Centrality and Focused Current Flow Betweenness Centrality. Note that the proposed semantic expansion achieves better precision results in presence of input

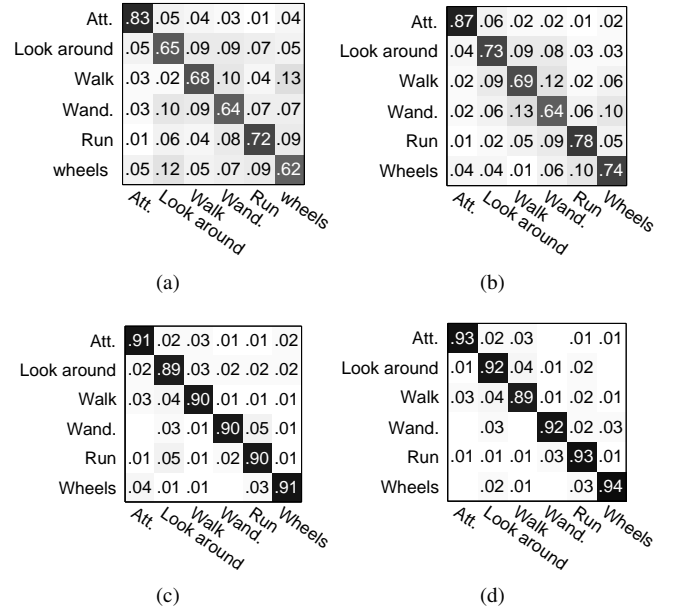


Fig. 6: Classification accuracy using different techniques: a) Lu et al. [1] b) Varini et al. [28] c) Poleg et al. [3] d) our approach.

terms linked to specific concepts related to the geographic location of the video. Considering, for example, the input user preferences “Sport, Cars” and the location “Modena”, we note the expansion includes several related terms strongly linked to the city such as “Ferrari”, “Maserati” etc. On the other hand the proposed solution obtains less accurate expansion in presence of term that are not strictly related to a specific location: for example the user preference “Fashion, Clothing, shopping” and the location “Modena”. Furthermore, the Table presents the comparison results using a node expansion of two and three hops (see III-B). In most of the cases the two hops expansion achieves a better performance in both precision and centrality metrics. For instance, in the user preference “Architecture, Renaissance”, location “Ferrara”, the three hops expansion adds some terms less related, such as “Music” and “Abstract expressionism”. Indeed, as the number of related nodes rears, the precision performance decreases from 91% to 72% and the value of centrality metric decreases by 23%, since the communities are enormously inflated and less focused.

TABLE II: Comparison of shots classification F1-score using the 3 aforementioned approaches.

Method	F1-score
No Preprocessing (NP)	39.75
Semantic Expansion (SE)	42.91
Visual image filtering (VIF)	46.7
Semantic Expansion plus Metadata and visual image filtering (SE-VIF)	55.75

Once the expanded user preferences is obtained we build a set of visual semantic classifiers. To analyze the effectiveness of our solution to detect user preferences into videos we annotate the full dataset (48 videos) according to a set of different user inputs. Table II shows the F1-Score results in the

¹<http://imagelab.ing.unimore.it/egosummarization>

TABLE III: Results of semantic expansion of different user preferences. Expanded preferences obtained using three hops only show the addition terms with respect to two hops results.

User preference	Expanded Preferences	Location	Hops	Precision	Avg. Node Connectivity	Avg.Focused Information Centrality
Church, Romanesque	Romanesque, Architecture, Cathedral, Church, Cathedral Torre Civica and Piazza Grande, Gothic, Lanfranco, Painting, Reliefs, Saint Geminianus, Sculpture, Willigelmus	Modena	2	92	2.25	0.635
Church, Romanesque	Romanesque, Medieval, Catholic, Basilica, Castle, Torre della Ghirlandina, Thomas of Britain, Tristan	Modena	3		1.98	0.462
Architecture, Renaissance	Art, Renaissance, San Pietro, Church, Architecture, Painting,	Modena	2	91	1.81	0.514
Architecture, Renaissance	Art, Renaissance, Catholic, Galleria Estense, Music Tintoretto, Veronese, Correggio, Dell'Abate	Modena	3	72	1.83	0.519
Architecture, Baroque	Baroque, San Agostino, San Barnaba, San Carlo, Chiesa del Voto	Modena	2	95	1.89	0.457
Architecture, Baroque	Baroque, Catholic, Basilica, Reni, Lana, Guercino, Velasquez, Bernini	Modena	3	97	1.88	0.523
Sport, Cars	Sportcars, Enzo Ferrari, Ferrari, Maserati, Lamborghini, de Tomasi, Pagani, Automobile, Manufacturer, Museum	Modena	2	99	1.75	0.467
Sport, Cars	Sportcars, Racecar Driver, Race, Rally, Team	Modena	3	90	1.74	0.433
Fashion, clothing, Shopping	Fashion, Clothing, Shopping, Stores, Online Shopping, Window Shopping, Mall	Modena	2	86	1.83	0.423
Fashion, Clothing, Shopping	Fashion, Clothing, Shopping, Industry, Activity, Retail	Modena	3	62	1.83	0.396
Reneissance, Architecture	Baroque, Architecture, Bentivoglio, Certosa, Church of Ges, Palazzo dei Diamanti, Giulio d'Este, Oratorio, Palace, Schifanoia, San Francesco,	Ferrara	2	100	2.96	0.536
Reneissance, Architecture	Baroque, Music, Bugnato, Garofalo, Mantegna, Gonzaga, Sigismondo d'Este, Abstract Espressionism	Ferrara	3	58	1.30	0.404
Medieval, Castle	Medieval, Bacchanalia, Castle, Chapel, Court, Dawn, Dungeons, Lions, Loggia, Tower,	Ferrara	2	100	2.96	0.596
Medieval, Castle	Medieval, Castle, Bacchus, Calvinism, Diana, Endymion, Evangelists, Pan	Ferrara	3	43	1.12	0.288

full dataset of four different strategies to grab online images for building visual classifiers:

- *No Preprocessing step*: we simply extract images from Visual Online Repositories (Flickr, Google) using the input user preferences.

- *Semantic Expansion* we download images from Visual Online Repositories using the terms of the expanded preferences.

- *Visual Image Filtering*: we extract images from Visual Online Repositories using the input user preferences. Images are then filtered considering jointly visual and semantic features extracted by their metadata.

- *Semantic Expansion plus metadata and visual image filtering (the proposed solution)*: we extract images from Visual Online Repositories using the expanded set of the user preferences. Images are further filtered considering jointly visual and semantic features extracted by their metadata.

We note that both semantic expansion and visual image filtering strategies, taken alone, achieve better classification performance with respect to the “No preprocessing” baseline procedure. Moreover their combination, “Visual Preprocessing plus Semantic Preprocessing” strategy, further improves the

user preference detection in the videos. In particular, Fig. 7 shows eight examples of classification comparison between the four approaches in term of F1-Score.

We observe that user preferences strongly related to cultural heritage sites of the city take a remarkable advantage. For instance, in Bologna, characterized by the presence of a large number of medieval, civil and religious, architecture buildings, user preferences such as “Medieval, Architecture” or “Church, Cathedral” can be easily expanded semantically without introducing noisy terms. Therefore, the extracted images from online repositories are more coherent to the specific visit (Fig 7 - videos: 1,2). Similarly, this effect can be noted in the other art cites; for example see in Ferrara and Ravenna (Fig 7 - videos: 3,4 and 7,8), respectively with “Reneissance, Architecture”, “Medieval, Castle” and “Dante, Medieval Architecture” (Video 4, 7, 8). On the other hand, for example, in the “shopping” case the expansion procedure alone introduces hypernymy terms that decreases its specificity. However, the visual filtering mitigates this effect, since it performs a clustering based on visual and semantic features and removes outliers items. It is also interesting to notice

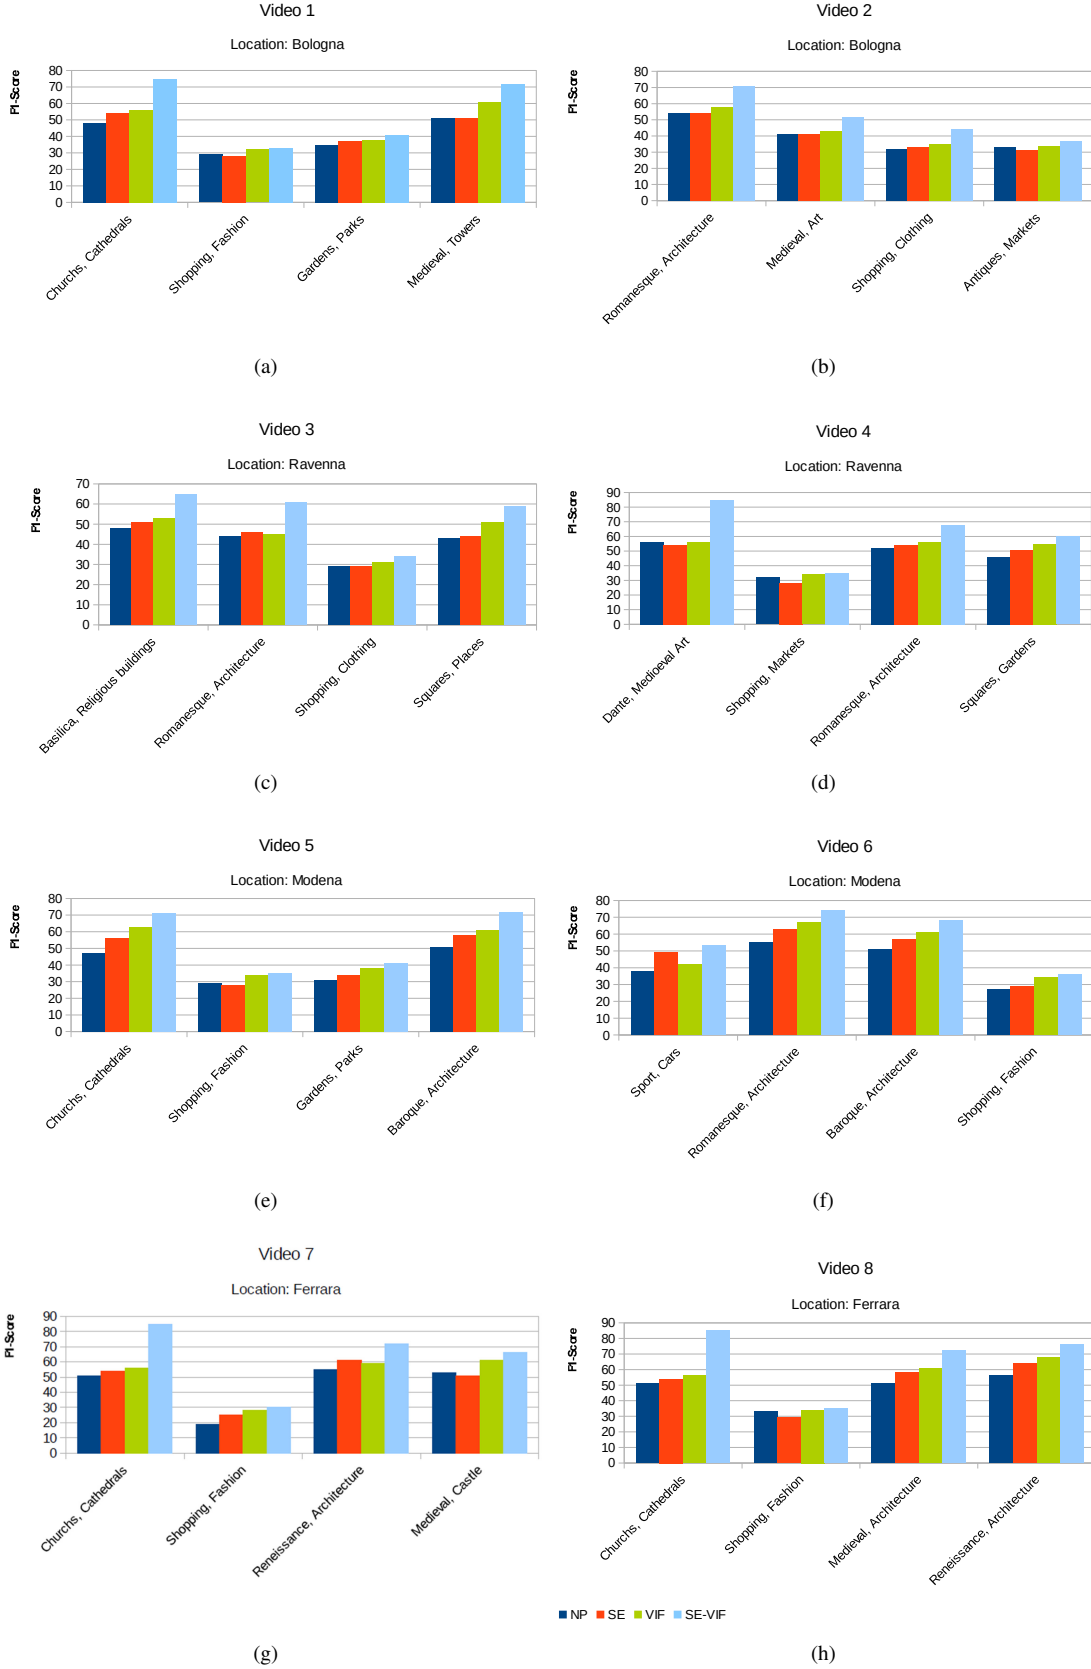


Fig. 7: For each Video the classification comparison between classifiers trained with “NP”=“No preprocessing”, “SE”=“Semantic Expansion”, “VIF”=“Visual Image Filtering” and “SE-VIF”=“Semantic Expansion plus Metadata and Visual Image Filtering”, as defined in Section IV-B is shown.




Fig. 8: “Blind taste test” user evaluation.

that for the user preference “Sport, Cars” in location Modena (Video 6), “Semantic Expansion” alone remarkably improves the user preference detection in the videos, while the “Visual Image Filtering” alone reduces this improvement. This is presumably due to the fact that semantic expansion introduces highly specifically related terms, while visual filtering on images extracted without semantic expansion suffers from the large heterogeneity introduced by showroom sports cars images not embedded in real environments.

C. Experiments on Narrativity

To evaluate the proposed narrative assessment strategy we define two Baselines. The first baseline (called “Uniform”) is obtained with a uniform sampling, i.e. given a video of length L , a desired length l and a minimum number of frames f for a ego-shot, an uniform selection of frame sequence at equal temporal length is made. For each selected frame the narrativity score is set to 1, otherwise 0. The second baseline (called “DBSCAN”) is obtained using the DBScan agglomerative clustering over dense layer pre-trained $C3D$ net [46] features extracted from ego-shots, identified by using behavior pattern analysis. In particular, this network is trained from scratch to extract a visual feature by using a subset of 2000 shots of our dataset. Given a desired length l , we automatically fix the number of ego-shots to arrange the final video. We choose, first, ego-shots belonging to larger clusters (in particular the nearest to the pseudo-centroids). As before, we assign score equal to one to the selected shots (zero otherwise).

To analyze our methodology, a subset of videos is annotated identifying ego-shots that are important milestone in the narrative flow. Table IV presents experimental results in terms of Precision, recall and F1-Score.

As can be seen our method outperforms largely both the baselines. Uniform baseline performs indeed better than DBSCAN baseline, presumably due to the chronological smoothness naturally achieved in uniform approach, while DBSCAN makes no assumption on temporal flowing of the scenes. In particular, it emerges that a relevant fraction of uninforma-

TABLE IV: Comparison results of Narrativity assessment in terms of F1-Score.

Attribute	Precision	Recall	F1-score
Uniform Baseline	23.95	38.80	29.62
DBSCAN Baseline	26.34	30.22	28.15
Our Narrative Approach	85.24	79.56	82.30

tive shots, containing accidental captures of walls, cobbled pavings, asphalt roads, grids yards, crowds, traffic and busy streets, is filtered by means of informativeness assessment; lower but significant was the recall for highly informative shots.

Colorfulness classification helped to filter a large number of visually dull scenes, as the ones in which the viewer was staring at phone, timetables, road signs, walls.

D. Experiments on Summarization

Finally, to evaluate the effectiveness of our solution to obtain a summarized videos based on user preferences, we perform a “blind taste test” on a group of twelve users of different sexes and ages (15-40 years old). For each video a participant has to choose which summary best meets user preferences among our proposed solution, from now on addressed as “User Preferences” (UP), and three different approaches.

The three competitors consist of the upper-cited “Uniform” Baseline and “DBSCAN” Baseline (described in Section IV-C), plus an approach that we denominate “Different user Preferences” (DUP). “Different User Preferences” (DUP) baseline is obtained by using our strategy, UP, but with different input terms.

We first show to the test group a browsable sped-up version of the entire original videos. Afterwards, for each original video, we show them the four summaries. We do not reveal the order as it is obtained randomly. After viewing all of them, each participant is asked to score summaries w.r.t. representativeness of user preference (“Representativeness Score”) and narrativity (“Narrativity Score”), assigning an exclusive rate between 0 and 3 for each assessed attribute. As can be seen, this test shows that 85% of the comparisons assigns

a higher score to summaries obtained with our approach in both evaluation criteria with respect to both baselines and “Different user Preferences”. Nevertheless it is interesting to notice that “Uniform Baseline” performs generally better than DBscan, as it renders more smoothly the chronological storyline. This result is more marked in narrativity assessment. In fact “DBscan”, in most of the cases, creates a very jumpy and discontinuous summary, as lacking temporal smoothness. It is also interesting to notice that “Different user Preferences” approach performs also generally better than the baselines. This is presumably due to the fact that, even if it is less or not at all focused on the topics of interest for the user, it provides both temporal smoothness, thanks to the narrativity component, and shakiness reduction due to the camera wearer’s behavior assessment.

V. CONCLUSIONS

We proposed a new method for personalized video summarization in Cultural Heritage Domain where we formulated the problem in terms of behavior, semantic relatedness to the expressed user preferences and narrativity. Behavior is learned from visual apparent motion features and real motion features measured by GPS sensors. Relatedness to user preference is assessed by means of on-the-fly created semantic classifiers. Narrativity is measured according to shots importance and relevance assessed through a Personalized Page Rank approach, with learned model parameters. The experiments show interesting results since summaries produced with our method have been voted as the most relevant in terms of fidelity to user preferences and chronological visual smoothness.

REFERENCES

- [1] Z. Lu and K. Grauman, “Story-driven summarization for egocentric video,” in *In Proc. of CVPR*, 2013.
- [2] Y. Poleg, T. Halperin, C. Arora, and S. Peleg, “Egosampling: Fast-forward and stereo for egocentric videos,” in *In Proc. of CVPR*, 2015.
- [3] Y. Poleg, A. Ephrat, S. Peleg, and C. Arora, “Compact cnn for indexing egocentric videos,” *arXiv preprint arXiv:1504.07469v1*, 2015.
- [4] Y. J. Lee, J. Ghosh, and K. Grauman, “Discovering important people and objects for egocentric video summarization,” in *In Proc. of CVPR*, 2012.
- [5] M. Gygli, H. Grabner, R. Hayko, and L. Van Gool, “Creating summaries from user videos,” in *In proc. of ECCV*, 2014.
- [6] B. Zhao and E. P. Xing, “Quasi real-time summarization for consumer videos,” in *In Proc. of CVPR*, 2014.
- [7] H. W. Ng, Y. Sawahata, and K. Aizawa, “Summarization of wearable videos using support vector machine,” in *In Proc. of IEEE International Conference on Multimedia and Expo*, 2002.
- [8] J. Xu, L. Mukherjee, Y. Li, J. Warner, J. M. Reh, and V. Singh, “Gaze-enabled egocentric video summarization via constrained submodular maximization,” in *In Proc. of CVPR*, 2015.
- [9] S. Swetha, A. Mishra, G. M. Hegde, and C. Jawahar, “Efficient object annotation for surveillance and automotive applications,” in *In Proc. of WACVW*, 2016.
- [10] T. Shu, D. Xie, B. Rothrock, S. Todorovic, and S.-C. Zhu, “Joint inference of groups, events and human roles in aerial videos,” in *In Proc. of CVPR*, 2015.
- [11] B. Li and M. I. Sezan, “Event detection and summarization in sports video,” in *Content-Based Access of Image and Video Libraries, 2001.(CBAIVL 2001). IEEE Workshop on*. IEEE, 2001, pp. 132–138.
- [12] B. Li, H. Pan, and I. Sezan, “A general framework for sports video summarization with its application to soccer,” in *Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP’03). 2003 IEEE International Conference on*, vol. 3. IEEE, 2003, pp. III–169.
- [13] F. Chen and C. De Vleeschouwer, “Formulating team-sport video summarization as a resource allocation problem,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 21, no. 2, pp. 193–205, 2011.
- [14] Z. Xiong, R. Radhakrishnan, A. Divakaran, Y. Rui, and T. S. Huang, *A unified framework for video summarization, browsing & retrieval: with applications to consumer and surveillance video*. Academic Press, 2006.
- [15] Z. Ji, Y. Su, R. Qian, and J. Ma, “Surveillance video summarization based on moving object detection and trajectory extraction,” in *Signal Processing Systems (ICSPS), 2010 2nd International Conference on*, vol. 2. IEEE, 2010, pp. V2–250.
- [16] Y. Peng and C.-W. Ngo, “Clip-based similarity measure for query-dependent clip retrieval and video summarization,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 16, no. 5, pp. 612–627, 2006.
- [17] M. Wang, R. Hong, G. Li, Z.-J. Zha, S. Yan, and T.-S. Chua, “Event driven web video summarization by tag localization and key-shot identification,” *IEEE Transactions on Multimedia*, vol. 14, no. 4, pp. 975–985, 2012.
- [18] T. C. Del Molino, Ana Garcia, J.-H. Lim, and A.-H. Tan, “Summarization of egocentric videos: A comprehensive survey,” *IEEE Transactions on Human-Machine Systems*, pp. 1–12, 2016.
- [19] M. Bolanos, M. Dimiccoli, and P. Radeva, “Towards storytelling from visual lifelogging: An overview,” *arXiv preprint arXiv:1507.06120*, 2015.
- [20] M. Gygli, H. Grabner, and L. Van Gool, “Video summarization by learning submodular mixtures of objectives,” in *In Proc. of CVPR*, 2015.
- [21] A. Lidon, M. Bolaños, M. Dimiccoli, P. Radeva, M. Garolera, and X. Giró-i Nieto, “Semantic summarization of egocentric photo stream events,” *arXiv preprint arXiv:1511.00438*, 2015.
- [22] N. Joshi, W. Kienzle, M. Toelle, M. Uyttendaele, and M. F. Cohen, “Real-time hyperlapse creation via optimal frame selection,” *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, p. 63, 2015.
- [23] M. Okamoto and K. Yanai, “Summarization of egocentric moving videos for generating walking route guidance,” in *Pacific-Rim Symposium on Image and Video Technology*. Springer, 2013, pp. 431–442.
- [24] Y.-L. Lin, V. I. Morariu, and W. Hsu, “Summarizing while recording: Context-based highlight detection for egocentric videos,” in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2015, pp. 51–59.
- [25] T. Yao, T. Mei, and Y. Rui, “Highlight detection with pairwise deep ranking for first-person video summarization.”
- [26] K. Zhang, W.-L. Chao, F. Sha, and K. Grauman, “Summary transfer: Exemplar-based subset selection for video summarization,” *arXiv preprint arXiv:1603.03369*, 2016.
- [27] H. W. Ng, Y. Sawahata, and K. Aizawa, “Summarization of wearable videos using support vector machine,” in *In Proc. of ICME*, 2002.
- [28] P. Varini, G. Serra, and R. Cucchiara, “Personalized egocentric video summarization,” in *In Proc. of ACM*, 2015.
- [29] J. M. Henderson, “Regarding scenes,” *Current Directions in Psychological Science*, vol. 16, no. 4, pp. 219–222, 2007.
- [30] G. Farneback, “Two-frame motion estimation based on polynomial expansion,” in *In Proc. of the 13th Scandinavian Conference on Image Analysis*, 2003.
- [31] F. Crété-Roffet, T. Dolmiere, P. Ladret, M. Nicolas *et al.*, “The blur effect: Perception and estimation with a new no-reference perceptual blur metric,” in *In proc. of SPIE*, 2007.
- [32] N. Cristianini and J. Shawe-Taylor, *An introduction to support vector machines and other kernel-based learning methods*. Cambridge university press, 2000.
- [33] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives, “Dbpedia: A nucleus for a web of open data,” in *The semantic web*. Springer, 2007, pp. 722–735.
- [34] E. W. Dijkstra, “A note on two problems in connexion with graphs.”
- [35] M. E. Newman and M. Girvan, “Finding and evaluating community structure in networks,” *Physical review E*, vol. 69, no. 2, p. 026113, 2004.
- [36] U. Brandes and D. Fleischer, “Centrality measures based on current flow,” in *In Proc. of Annual Symposium on Theoretical Aspects of Computer Science*, 2005.
- [37] O. O. Beineke, L. and R. Pippert, “The average connectivity of a graph,” *Discrete mathematics*, vol. 3, pp. 31–45.
- [38] I. Hulpus, C. Hayes, M. Karnstedt, and D. Greene, “Unsupervised graph-based topic labelling using dbpedia,” in *In Proc. of ACM*, 2013.
- [39] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.

- [40] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [41] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *Kdd*, vol. 96, no. 34, 1996, pp. 226–231.
- [42] M. Gygli, H. Grabner, and L. Gool, "Video summarization by learning submodular mixtures of objectives," in *In Proc. of ECCV*, 2015.
- [43] S. D. Kamvar, T. H. Haveliwala, C. D. Manning, and G. H. Golub, "Extrapolation methods for accelerating pagerank computations," in *In Proc. of International conference on WWW*, 2003.
- [44] A. N. Langville, C. D. Meyer, and P. Fernández, "Googles pagerank and beyond: The science of search engine rankings," *The Mathematical Intelligencer*, vol. 30, no. 1, pp. 68–69, 2008.
- [45] D. A. Schult and P. Swart, "Exploring network structure, dynamics, and function using networkx," in *In Proc. of SciPy Conf.*, 2008.
- [46] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "C3d: generic features for video analysis," *CoRR, abs/1412.0767*, vol. 2, p. 7, 2014.
- [47] X. Lu, Z. Lin, H. Jin, J. Yang, and J. Z. Wang, "Rapid: Rating pictorial aesthetics using deep learning," in *ACM Multimedia*, 2014.
- [48] R. Collobert and J. Weston, "A unified architecture for natural language processing: Deep neural networks with multitask learning," in *In Proc. of ICML*, 2008.