

Semantic Summarization of Egocentric Photo Stream Events

Aniol Lidon
Universitat Politècnica de Catalunya
Barcelona, Catalonia/Spain

Marc Bolaños
Universitat de Barcelona
Barcelona, Spain
marc.bolanos@ub.edu

Mariella Dimiccoli
Universitat de Barcelona
Computer Vision Center
Barcelona, Spain
mariella.dimiccolig@cvc.uab.es

Petia Radeva
Universitat de Barcelona
Barcelona, Spain
petia.radeva@ub.edu

Maite Garolera
Consorci Sanitari de Terrassa
Terrassa, Spain
maite.garolera@cst.cat

Xavier Giro-i-Nieto
Universitat Politècnica de Catalunya
Barcelona, Catalonia/Spain
xavier.giro@upc.edu

ABSTRACT

With the rapid increase of users of wearable cameras in recent years and of the amount of data they produce, there is a strong need for automatic retrieval and summarization techniques. This work addresses the problem of automatically summarizing egocentric photo streams captured through a wearable camera by taking an image retrieval perspective. After removing non-informative images by a new CNN-based filter, images are ranked by relevance to ensure semantic diversity and finally re-ranked by a novelty criterion to reduce redundancy. To assess the results, a new evaluation metric is proposed which takes into account the non-uniqueness of the solution. Experimental results applied on a database of 7,110 images from 6 different subjects and evaluated by experts gave 95.74% of experts satisfaction and a Mean Opinion Score of 4.57 out of 5.0. Source code to reproduce this work is available at <https://github.com/imatge-upc/egocentric-2017-lta>.

ACM Reference format:

Aniol Lidon, Marc Bolaños, Mariella Dimiccoli, Petia Radeva, Maite Garolera, and Xavier Giro-i-Nieto. 2017. Semantic Summarization of Egocentric Photo Stream Events. In *Proceedings of LTA'17, Mountain View, CA, USA, October 23, 2017*, 9 pages.
<https://doi.org/10.1145/3133202.3133204>

1 INTRODUCTION

From smartphones to wearable devices, digital cameras are becoming ubiquitous. This process is being accompanied by the progressive reduction of digital storage cost, making it possible to collect large amounts of high-quality pictures in a easy and affordable way. This situation arises a number of natural questions: how to manage this large amount of pictures? Do we really need to store all of them? Summarization, the process of generating a proper, compact and meaningful representation of a given image collection through a subset of representative images, is crucial to help managing and

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

LTA'17, October 23, 2017, Mountain View, CA, USA

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5503-2/17/10...\$15.00

<https://doi.org/10.1145/3133202.3133204>



Figure 1: Example of temporal neighbouring images acquired by a wearable photo camera.

browsing efficiently large volumes of video content. Although summarization is not a new research topic in Computer Vision [13, 29], it is still a largely open problem.

Our goal in this paper is to address the summarization problem focusing on a particular scenario: the one of analyzing photo streams acquired through a wearable camera. Motivation for this work is given by the explosion of the number of wearable camera users in recent years and, consequently, by the huge amount of data they produce. To record daily experiences from an egocentric, first-person perspective is a trend that has been growing progressively since 1998, when Steve Mann proposed the WearCam [23]. In 2000, Mayol et al. proposed a necklace-like lifelogging device [24] and, in 2006, Microsoft Research started to commercialize the first egocentric lifelogging portable camera, the SenseCam, for research purposes [14].

Several authors [21, 27, 28, 31] have studied the benefits of lifelogging cues such as egocentric images to help people with dementia to enhance their memory or to help them remember about their forgotten past. Sellen et al. [31] showed that episodic details from a visual 'lifelog' can be presented to users as memory cues to assist them in remembering the details of their original experience. To support people with dementia, Piasek et al. [27] introduced the "SenseCam Therapy" as a therapeutic approach similar to the well established "Cognitive Stimulation Therapy" [33]. Participants were asked to wear SenseCam in order to collect images of events from their everyday lives, then images were reviewed with a trained therapist.

However, lifelogging technologies produce huge amounts of data (thousands of images per day) that should be reviewed by both patients and caregivers. To be efficient, lifelogging systems need to summarize the most relevant information in the images. On the other hand, in order to make possible the recording of images from the whole day, it is necessary to use wearable cameras with low temporal resolution (2 fpm). The peculiarity of these image collections is that, due to the low-temporal resolution of the camera

and to the free motion of its wearer, temporally adjacent images may be very different in appearance even if they belong to the same event and should therefore be grouped together. For instance, during a meeting, the people the camera wearer is interacting with and the objects around them may change their position and frequently appear occluded (see Fig. 1).

As a consequence of this, taking the semantics into account is crucial to summarize egocentric sequences acquired by a low temporal resolution camera. This paper proposes a method to summarize each of the events present in a daily egocentric sequence, aiming at preserving semantic information and diversity, while reducing the total number of images. The method consists of three major steps: first, non-informative images are removed; second, they are ranked by semantic relevance; and finally, a new re-rank is applied by enforcing diversity among the chosen subset of pictures. Our contributions can be summarized as follows:

- (1) Propose a CNN-based informativeness estimator for egocentric images.
- (2) Define a set of semantic relevance criteria for egocentric images.
- (3) Formulate the summarization task as a retrieval problem by combining informativeness, relevance and novelty criteria.
- (4) Define a soft metric to assess the novelty from partially annotated image datasets.
- (5) Our results have been validated by medical experts with the aim of being used in a cognitive training framework to reinforce the memory of patients with mild cognitive impairment.

The rest of the paper is organized as follows: the next section reviews the related work, section 3 details the proposed method, section 4 describes our experimental setup, section 5 presents the experimental results and finally, section 6 ends the paper with some concluding remarks. Source code to reproduce this work is available at <https://github.com/imatge-upc/egocentric-2017-lta>.

2 RELATED WORK

2.1 Summarization by key-frame selection in Lifelogging

Egocentric photo stream summarization has been traditionally formulated as the problem of grouping lifelog images into coherent collections (or events) by: first, extracting low-level spatio-temporal features and then, selecting the most representative image from each event. In this spirit, many authors proposed different strategies for temporal segmentation and key-frame selection. Doherty et al. [10] proposed a key-frame selection technique, which seeks to select the image with the highest 'quality' as key-frame by relying on five types of features: contrast, color variance, global sharpness, noise, saliency and external sensors data (accelerometers and light). In addition to image quality, Blighe et al. [2] considered image similarity to select a key-frame in an event. Basically, after applying an image quality filter that removes all poor quality images, they selected the image which has the highest average similarity to all other images in the event as the key-frame. In that case, similarity was measured relying on the distance of SIFT descriptors. In [18], the key-frame is selected by a nearest neighbour ratio strategy that

favors high quality images and, if the difference in quality is not large enough, favors images closer to the middle of the temporal segment. More recently, the authors in [3] proposed to use a Random Walk for selecting a single and most representative image for each temporal segment.

While these methods rely solely on low-level or mid-level features for the temporal segmentation and key-frame detection, a few recent works have introduced a higher semantic level in the selection process for video cameras. Although, in this cases, due to the higher temporal resolution of the camera (about 30fps), aiming at selecting subshots (short video sub-sequences) instead of unique key-frames. Lu and Grauman [22] and lately Ghosh et al. [11] suggested that video summarization should preserve the narrative character of a visual lifelog and, therefore, it should ideally be made of a coherent chain of video subshots in which each subshot influences the next through some subset of key visual objects. Following this idea, in [11, 22], first important people and objects are discovered based on their interaction time with the camera wearer and then, a subshot selection driven by key-object event occurrences is applied. Subshot selection is performed by incorporating into an objective function a term corresponding to the influence between subshots as well as image diversity. To model diversity, the authors relied on GIST descriptors and color histograms to model scenes and proposed a measure of diversity that is high when the scenes in sequential subshots are dissimilar to ensure visual uniqueness. Visual diversity in lifelog summaries is modeled in [1] through the concept of novelty, which the authors heuristically defined as the deviation from some standard background. According to this definition, novelty is detected based on the absence of a good registration, in terms of ego-motion and the environment, between a new sequence and stored reference sequences. More recently, Gong et al. [12] proposed the so called Sequential Determinantal Point Process (seqDPP) approach for video summarization, a probabilistic model with the ability to teach the system how to select informative and diverse subsets from human-created summaries, so as to best fit human-perception quality based on evaluation metrics. This work is an adaptation to sequences of [20], which is able to capture the strong dependency structures between items in sequences. It is worth to mention that all these works have been conceived to deal with video data, where, assuming that the temporal segmentation is good, temporal coherence and frame redundancy make the key-frame selection process easier.

2.2 Diversity and novelty in information retrieval

One of the first works that tried to approach the problem of obtaining a diverse set of elements was presented in 1998 by Carbonell & Goldstein [4]. Their proposal, which was applied to the context of text retrieval and summarization, aimed at obtaining results highly relevant for the query, but presenting a low redundancy. They define the marginal relevance as the linear combination of relevance and novelty, measured independently. They aimed at maximizing it iteratively, defining this way what they call the Maximal Marginal Relevance (MMR).

A similar formulation to MMR of the diversification problem was given more recently by Deselaers et al. [8] in the problem of image

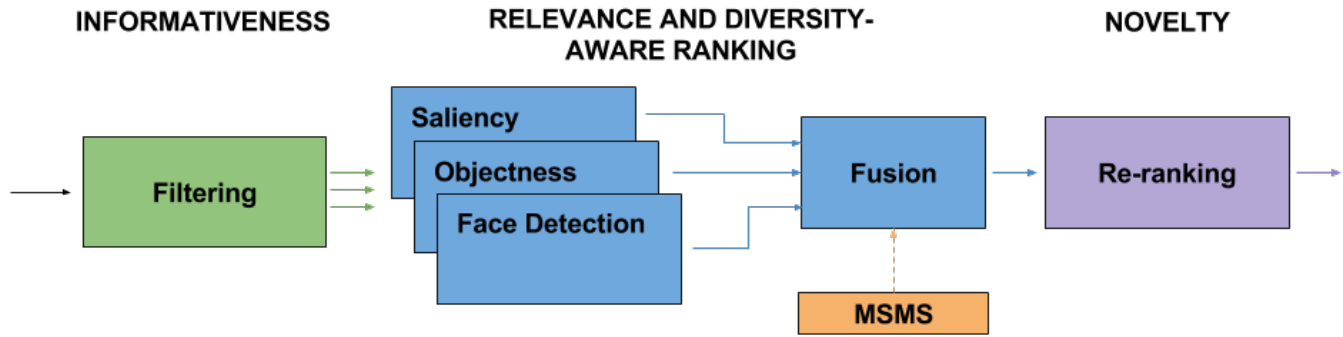


Figure 2: General scheme of our event summarization methodology.

retrieval. Likewise, they jointly optimized the relevance and the diversity of the query, although reformulating the general solution by incorporating dynamic programming techniques to the initial proposal.

Diversity in social image retrieval was one of the focus of the MediaEval 2013, 2014 and 2015 benchmarks and attracted the interest of many groups working in this area. Most participants developed diversification approaches that combined clustering with a key-frame selection strategy to extract representative images for each cluster. Spyromitos et al. proposed an MMR-based approach [34] that jointly considers relevance and diversity, but using a supervised classification model to obtain the relevance scores learned from the user feedback. The contribution from Dang-Nguyen et al. [6] was to filter out non-relevant images at the beginning of the process before applying diversity, hence simplifying it.

In [32], the authors presented a method for detecting and resolving the ambiguity of a query based on the textual features of the image collection. If a query has an ambiguous nature, this ambiguity should be reflected in the diversity of the result to increase user satisfaction. Leuken et al. [37] proposed to reduce the ambiguity of results provided by image search engines relying on textual descriptions by seeking for the visual diversification of image search results. This is achieved by clustering the retrieved images based on their visual similarity and by selecting a representative image for each cluster. In these works 'diversity' is aimed at addressing the 'ambiguity' of the textual descriptions (tags) in queries rather than avoiding redundancy in search results. The same terminology is used by [5], where the term 'novelty' is meant to address redundancy in the retrieved documents. In this work we will use the terms novelty and diversity as defined in [5].

3 METHODOLOGY

Inspired by the image retrieval work, we question which would be the most suitable diversity and relevance criteria applied on egocentric images that kept the narrative character of the visual lifelog. Taking into account that the images are acquired non-intentionally, it is important to disregard the non-informative images (see Section 3.1) and keep the minimal set that represents the visual event. In this section, we explain the four main steps applied to construct the final resulting summary (see Fig. 2):

(1) **Informativeness Filtering** (Section 3.1): egocentric images are acquired non-intentionally and therefore, many of them may be non-informative, capturing neither (or partially) objects nor people, or being blurred or dark. By means of a CNN-based informativeness filtering method, we discard most of the non-informative images from the egocentric event.

(2) **Relevance and Diversity-aware Ranking** (Section 3.2): an initial relevance image ranking is computed taking into account different criteria such as Saliency and Objectness for dealing with the ambiguity or under-specification of queries like *What is the user doing?*, or *Where is the user?*, and Face Detection for answering the question *With whom is the user most likely interacting?*

(3) **Novelty-based Re-ranking** (Section 3.3): a final re-ranking based on a novelty-maximization procedure is applied on the images already ranked by relevance. This step is crucial to select those images that represent the most varied set of concepts appearing avoiding redundancy without semantic loss.

(4) **Estimation of Fusion Weights with Mean Sum of Maximal Similarities** (Section 3.4): in this step, we define a novel soft metric, which we call Mean Sum of Maximal Similarities (MSMS), in order to define the priorities of the different relevance terms and construct the final summary.

3.1 Informativeness Filtering

Considering both the free motion of egocentric cameras and the non-intentionality of the pictures they take, several problems are inherent to them as over- or under-light exposure, blurriness, pictures of the sky or ground, pictures where possible objects of interest are badly centered in the image or even non-existent, etc. Such images are considered as *non-informative pictures*. A way to avoid a good amount of the undesired images in the final summary and also to boost the performance of the next steps in our methodology, would be being able to discriminate the *informative pictures*. Fig. 3 illustrates informative vs. non-informative photos taken by the wearable camera.

In order to learn a model of informative pictures, and taking into account the complexity of distinguishing visually whether an image is informative enough or not, we propose training a CNN for a binary problem. Thus, all the images with an undesired artifact (empty image, blurred image, image with small amount of information, image where something occludes most of the image




Figure 3: Examples of informative images (top) and non-informative images (bottom) belonging to the same events. Faces appearing in the images have been manually blurred for privacy concerns.



Figure 4: Images acquired by a lifelogging device, where objects of interest appear like: computer, mobile, coffee, hand, bicycle, person, face, flower, apple, etc.

region, sky, ground, ceiling, room wall, etc.) will be considered non-informative (label 0), and the rest (images with semantic content) will be considered informative (label 1). With this procedure, we will be able to extract an *informativeness score* for each image and filter the unuseful ones. In order to remove as many images as possible, but with the ultimate goal of having a very high recall (always try to keep any informative image for the next steps), we only discard the images with *informativenessScore* < 0.025 , which correspond to the ones that are considered non-informative for certain by the CNN (see experimental results in section 5.1).

The network training for the binary class distinction is performed by fine-tuning the CaffeNet [19] pre-trained on the ImageNet [7] dataset, provided in the software Caffe [17].

3.2 Relevance and Diversity-aware Ranking

Once non-informative images are discarded, we proceed to rank the remaining images by considering a relevance criteria. Our solution formulates the summarization problem in similar terms as in information retrieval. A ranked list of event frames is generated in such a way that a summary of T images directly corresponds to a truncation of the ranked list of N elements at its T -th position. Classic retrieval problems build their ranked lists as a response to a user query, in the summarization of events, this query would correspond to: *Select T images to describe the event depicted into these N frames.* This query is highly ambiguous as an event can be described from different perspectives. For example, an egocentric visual summary of an event may respond to multiple intentions, such as: *Where is the user? What activity is the user performing? With whom is the user interacting?* We hypothesize that the relevance of each image with respect to these questions can be estimated with computer vision tools for saliency prediction, detection of objects and detection of faces.

3.2.1 Saliency Prediction for image relevance. We assume that images with more salient content are relevant and should have a higher probability to be included in the summary. Visual saliency

can be triggered by a broad range of reasons, such as objects, people or characteristic features appearing in the picture. Many computer vision algorithms try to estimate the fixation points of the human eyes in a scene by means of *saliency maps*. These are heat maps of the same size of the image, whose higher values correspond to the image locations with a higher probability of capturing the human visual attention.

In this work, we compute the saliency maps with SalNet [26], an end-to-end convolutional network for saliency prediction. We adopt the overall sum of the values in the saliency map as a quantitative estimator of the image relevance.

3.2.2 Object Detection for image relevance. We assume that those images containing objects are relevant since these objects likely correspond to the ones the user is interacting with. These objects would address the question of *What is the user doing?* (see Fig. 4). In addition, introducing a semantic interpretation of the scene, targets the summary from a higher abstraction level than saliency maps.

In our work, we used the off-the-shelf tool *Large Scale Detection through Adaptation (LSDA)* presented in [15]. This object detector is based on a CNN fine-tuned for local scale and provides a semantic label and a confidence score in the localization of the objects in the scene.

This tool allows to estimate the relevance of each frame by summing the detection scores of all objects in the picture, so that the frames with higher confidence detection will be considered as more relevant.

3.2.3 Face Detection for image relevance. We assume that images containing people are also relevant, since these people likely correspond to the ones the user is interacting with, and they would be useful to answer the question *With whom is the user interacting?* Therefore, a face detector complements the object detector into providing a cue for the user social interactions during the event.

In our solution, we adopted the off-the-shelf face detector by Zhu et al. [39], which provides a confidence for each of the detected faces. The relevance of each frame in terms of social interaction was estimated by summing the confidence scores of the face detectors. In the particular implementation of [39], detection scores may also be described with negative values, so we actually used these scores in an exponential sum, which conveniently deals with the negative scores as well as encourages the selection of frames with multiple detected faces.

3.2.4 Diversity-aware Ranking. The three criteria used for the relevance detailed above allow to cope with the ambiguity of the query, since they estimate the relevance from three different perspectives. The relevance scores computed for each case are then used to build three ranked lists that will be combined into a single one. This combination is based on generating a set of normalized scores based simply on the position of the frames. Normalized scores $r_k(x)$ are linearly distributed from the top (1 for most relevant) to the bottom (0 for non relevant) of the ranked list as follows:

$$r_k(x) = \frac{M - R_k(x)}{M - 1},$$

where $R_k(x) = 1, \dots, M$ is the ranking position associated with image x according to each relevance criterion $k \in \{1, 2, 3\}$ and M

is the number of informative frames in the event, being $M < N$ and N the total number of images in the event. The standard score normalization [25, 30] with the min and max scores was also tested giving similar results, so the rank-based normalization was adopted to save computational resources.

3.3 Novelty-based Re-ranking

The relevance and diversity-aware ranked list sorts the informative images combining different criteria for relevance, but does not explicitly cope with the redundancy in the information. Further processing is necessary to maximize the *novelty* provided by each image with respect to the rest in the summary. In information retrieval, *novelty* is defined as a quality of a system that avoids redundancy [5]. In our approach for building a summary by truncating a ranked list of images, introducing novelty implies that each image in the list should differ as much as possible from its predecessors.

Our approach adopts the *greedy* selection algorithm presented by Deselaers et. al [9] to re-rank the fused list based on novelty. The goal of our summarization algorithm is to analyze the input set of informative and ranked images $\mathcal{X} = \{x_1, \dots, x_M\}$ to iteratively build another set with minimal redundancy, say $\mathcal{Y}^T = \{x_{y_1}, \dots, x_{y_T}\}$, where $T \leq M$. Our approach starts by selecting the top ranked image x_1 in the diversity-aware ranked list as the first element of \mathcal{Y} , that is $\mathcal{Y}^1 = \{x_{y_1}\}$. The *novelty* of each candidate image, x^* to be added to the summary at iteration $t = 2, \dots, T$ is defined as:

$$n(x^*, \mathcal{Y}^t) = 1 - s(x^*, \mathcal{Y}^t) = 1 - \max_{x_{y_j} \in \mathcal{Y}^t} s(x^*, x_{y_j}) \quad (1)$$

$s(x^*, x_{y_j})$ is a normalized similarity measure between the candidate x^* and image x_{y_j} . In this way, the more different a new image is with respect to the ones in \mathcal{Y}^t , the higher its novelty is.

In our work, the similarity $s(x^*, x_{y_j})$ is based on visual appearance. Each image is described by a feature vector corresponding to the seventh fully-connected layer of the CaffeNet convolutional neural network [17], which was also used as initialization for the informativeness (see section 3.1), and whose architecture was inspired by AlexNet [19] and trained with the ImageNet dataset [7]. The similarity score is computed with the Euclidean distance of the feature vectors.

The summary of the event is built by combining the relevance and the novelty of the candidate frames, x^* . A greedy selection algorithm iteratively chooses the next image in the summary as:

$$\begin{aligned} x_{y_{t+1}} &= \arg \max_{x^* \in \mathcal{X} \setminus \mathcal{Y}^t} (r(x^*) + n(x^*, \mathcal{Y}^t)), \\ \mathcal{Y}^{t+1} &= \mathcal{Y}^t \cup \{x_{y_{t+1}}\} \end{aligned} \quad (2)$$

that is, the image with the highest sum of relevance and novelty. We detail the computation of $r(x^*)$ in Equation (4).

3.4 Estimation of Fusion Weights with Mean Sum of Maximal Similarities

Different relevance criteria would have different importance for the final ranking. To this aim, we propose a novel approach to estimate the priorities of each criterion and fuse them, based on comparing the similarity of the images from a validation set of

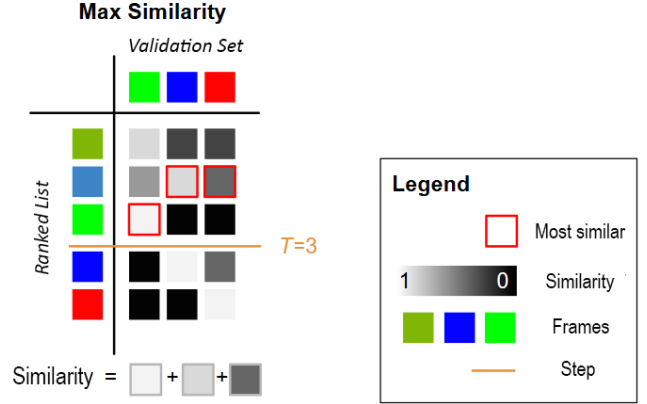


Figure 5: Visual representation of the Sum of Max Similarities when comparing the first $T = 3$ images from the ranked list with the $P = 3$ red, green and blue images from the validation set. The similarity between images from the validation set and images from the ranked list is represented as a $T \times P$ matrix of gray level squares, where each square (i, j) of the matrix represents the similarity between the image i from the ranked list and the image j from the validation set. High intensity gray level values indicate high similarity.

P elements $\mathcal{V} = \{x_{v_1}, \dots, x_{v_P}\}$ with a summary of T elements $\mathcal{Y}^T = \{x_{y_1}, \dots, x_{y_T}\}$.

The *Sum of Maximal Similarities* (SMS) of \mathcal{V} with respect to \mathcal{Y}^T is defined as:

$$SMS(\mathcal{V}, \mathcal{Y}^T) = \frac{1}{P} \sum_{i=1}^P s(x_{v_i}, \mathcal{Y}^T) \quad (3)$$

where $s(x_{v_i}, \mathcal{Y}^T)$ is the similarity of image x_{v_i} from the validation set with respect to \mathcal{Y}^T . Following this metric, the more similar is the validation set \mathcal{V} to the selected images \mathcal{Y}^T , the highest the average similarity.

The main advantage of our soft metric presented in Equation (3) is that automatic summaries, which are very similar to the validation images, although do not coincide, still will be assessed as a valid solution. This feature of our soft metric is specially important when working with sequences of egocentric photo streams, that usually contain a high amount of redundancy.

Our summaries are built as a truncation of a ranked list, so the value of SMS depends on T . Fig. 5 shows a schematic example, where a validation set composed of $P = 3$ images (represented by three colors: green, blue and red) is compared with an event of $M = 5$ images, which have been previously filtered and ranked based on our diversity and novelty-aware criterion. In this example, a summary with $T = 3$ images is considered. Let us imagine that the green image from \mathcal{V} is matched with the third one in the ranked list, while the blue and red images are matched with the second one.

Notice that, as in our set up the validation set \mathcal{V} is always a subset of the input set \mathcal{X} , hence, the final average similarity (i.e. considering the whole sequence, $T = M$ images) for \mathcal{Y}^M will always




Figure 6: MSMS curves and AUC for each relevance criterion used separately.

correspond to one. Actually, in the example of Fig. 5, SMS must reach a value of one when all the ranked list is considered, that is, when $T = M$.

Let us consider the evolution of the \overline{SMS} of summary \mathcal{Y} as the parameter t grows ($t = 1, 2, \dots, M$) defined as:

$$\overline{SMS}(\mathcal{V}, \mathcal{Y}) = \{SMS(\mathcal{V}, \mathcal{Y}^1), \dots, SMS(\mathcal{V}, \mathcal{Y}^M)\}$$

We construct the \overline{SMS} curves for all the validation set, interpolate them, normalize them with respect to the length M of each sequence and get the average of the curves, one curve for each event in the validation set. The resulting curve represents the *Mean Sum of Maximal Similarities (MSMS)*.

Fig. 6 shows three examples of MSMS curves. The curve illustrates the evolution of the MSMS as a function of the percentage of images covered by the visual summary. This curve is defined over the X -axis representing the proportion of event images represented in the top t items in the list, that is, plotting the $MSMS$ over t/M , where M is the amount of already filtered images in the event. In this way, the curves of events of different lengths M can be compared on the same plot. The best MSMS curves are those which reach higher values with the minimal amount of images in the summary.

Note that the MSMS can be computed for the three relevance criteria that can give as a quality of performance of each of them. Hence, we introduce a final refinement in the fusion of the ranked lists by estimating the confidence of each of the three relevance criteria and using it to weight the corresponding relevance terms. Thus, we leverage the contribution of each of the three relevance criteria based on their stand-alone performance (see Fig. 6). This weight corresponds to the normalized Area Under the Curve (AUC) of the MSMS measure.

As a result, the fusion weight $w(k)$ for the relevance criterion k is estimated as:

$$w(k) = \frac{AUC(k)}{\sum_{i=1}^3 AUC(i)}.$$

Finally, the three normalized relevance scores associated to each frame are aggregated with a weighted sum to obtain their fused score $r(x)$, as described by the following equation:

$$r(x) = \sum_{k=1}^3 w(k) r_k(x), \quad (4)$$

and the final set of frames is obtained according to the updated relevance and novelty criteria according to Equation (2).

4 EXPERIMENTAL SETUP

4.1 Dataset

The dataset used for validation of our method was acquired with the wearable camera Narrative Clip (www.getnarrative.com), which takes a picture every 30 seconds (2 fpm). It is composed of 10 day lifelogs from 5 different subjects, with a total of 7,110 images¹. Each day has been segmented in between 10 to 25 events manually, although any automatic segmentation method can be used (e.g. [36]). The event segmentation separated the pictures in a set of ordered and relevant semantic events or segments. In this paper, we use this segmentation as starting point for the semantic summarization.

In order to apply a quantitative evaluation of our method, each day lifelog was annotated at two levels by psychologists in the following way:

(1) **GT Level 1 - Informativeness:** positive or negative label depending on whether the image is considered informative or non-informative (see section 3.1). The proportion of images labeled as informative is 61.22% of the complete dataset.

GT Level 2 - Grouping of similar images: all highly similar informative images that belong to the same event are grouped together (see section 3.2.4). This distinction, resembling the one used in *MediaEval 2014 Retrieving diverse social images challenge* [16], intends to provide a way of measuring how many different clusters/groups from the ground truth are represented among the results in the final diversity selection.

4.2 Semantic Assessment of Summaries

The assessment of visual summaries is a challenging task due to the rich semantic content of the images and the ambiguity in the evaluation criteria. In addition, human annotation is an expensive resource, which becomes dramatically scarce, when dealing with tasks that require expert annotations in the domain. In our case, the ground truth annotations have been performed by psychologists who defined the event's summaries from an human-centered point of view.

We addressed the challenge of summary assessment as follows: starting from the daily events defined by the psychologists, the automatic summaries were built for each of the events separately by using the best configuration of our system. Later, the results were presented in a blind taste test to the expert annotators.

4.2.1 Validation based on MSMS. The classical methods to evaluate summaries cannot be applied in our setup, because they require the annotation of the full dataset. In a traditional case, each document in the ground truth is labeled as relevant or non-relevant for the summary and, in some cases, also clustered in groups of redundant items that cover the same sub-topic. Such annotations allow the definition of metrics like *Precision* [4] for relevance and

¹The link to the dataset and its ground truth are prepared to be made public domain when the article is published.

Cluster Recall (or subtopic-recall) [38] for diversity and/or novelty. However, in other cases (like ours), only a small portion of the dataset is annotated.

Given this limitation, following the framework of Subsection 3.3, we propose an evaluation approach based on the similarity of the images in the ground truth set $\mathcal{G} = \{x_{g_1}, \dots, x_{g_P}\}$, when compared to the automatic summary of T elements, $\mathcal{Y}^T = \{x_{y_1}, \dots, x_{y_T}\}$. In this case, we apply the SMS, $SMS(\mathcal{G}, \mathcal{Y}^T)$ of ground truth \mathcal{G} with respect to the extracted summary, \mathcal{Y}^T . As before, we obtain the MSMS as the average value, when considering all events in the dataset of the ground truth, which is finally plotted in the Y -axis of our evaluation scores. Finally, the AUC is computed to obtain a quantitative measure of each configuration for the summary.

4.2.2 Blind taste test. The resulting summaries are evaluated with a *blind taste test* by the team of experts who generated the ground truth summaries. This methodology, previously used in [3, 22], shows different summaries from the same event to the experts, so they can rate them from a comparative perspective. By randomly changing the position of the different techniques at each event, graders are unaware of what technique was used to generate each of the summaries, guaranteeing this way that their judgments are not biased towards any of the algorithm components.

The amount of images included in each summary corresponds to the number of frames selected by the experts, when building the ground truth summaries, that is, $P = T$. In this way, the summary shown to the graders is a truncation of the final ranked list of images. In our case, the experts decided the length of the automatic summary to be a percentage of the length of the whole event. Note that developing an automatic system to establish the length of the summary is out of the scope of this paper.

Additionally, the selected images are sorted in the summary according to their temporal time stamps. This final sorting helps the user to reconstruct the story between the images to have a better understanding of the event.

An online platform was developed for the experts to evaluate our results. In a first round of questionnaires, we wanted to compare the result of our diversity method using ImageNet or Places as a similarity criterion (see Equation (1)). For each event, the questionnaire first shows all the images belonging to the event, and then the user has to answer the questions *Is the summary representative of the event?* and *Which summary do you prefer?* for each similarity method (using ImageNet or Places).

On the second questionnaire the expert had to give a grade to each of the presented ranked summaries. The three summaries correspond respectively to: our approach, a uniform sampling of images as a lower-bound example, and the ground truth summaries constructed by the same experts two months earlier as an upper-bound of the performance score. The order in which the different summaries are shown is also randomly chosen for each event to avoid again any bias due to the sorting.

In this case, a comparison between the three types of summaries is obtained by asking the experts to grade each solution from 1 (worst) to 5 (best). This data collected allowed the computation of the *Mean Opinion Score (MOS)* for each configuration (see section 5.3).

5 RESULTS

Two types of experiments validate the presented results: one for the informativeness filtering and a second one using online questionnaires for the relevance, diversity and novelty on the final visual summaries. In both cases, the evaluation was based on the feedback provided by the experts.

5.1 Informativeness Filtering

To validate how successful is our algorithm to filter non-informative frames, we applied a 10-fold (a day set out) cross-validation. For all the experiments, we used the following parameters: $base_lr = 10^{-8}$, $lr_policy = "step"$, $gamma = 0.1$, $stepsize = 3000$, $blobs_lrf_{c6} = blobs_lrf_{conv5} \times 5$, $blobs_lrf_{c7} = blobs_lrf_{conv5} \times 8$ and $blobs_lrf_{c8} = blobs_lrf_{conv5} \times 10$. In Table 1, we present the validation accuracy obtained on each of the sets and the number of iterations applied to achieve the best result. After having the networks trained for each cross-validation, we evaluated the general average performance in terms of Accuracy, Precision, Recall and F-Measure for different informativeness score threshold values². As we can see, we are able to obtain the highest F-Measure value, when we filter images with an *informativenessScore* < 0.025 . In Fig. 7, we show the Precision-Recall curves for all thresholds, obtaining the best results with respect to the F-Measure of: $FMeasure = 0.881$, $Accuracy = 0.846$, $Precision = 0.84$ and $Recall = 0.926$ with *threshold* = 0.05.




Figure 7: Precision-Recall curve for different informativeness score threshold values (a small subset of them are shown in black dots) for all the sets.


5.2 Diversity

Fig. 8 illustrates qualitatively the differences obtained when introducing diversity to the ranked list. As we can see, when we introduce the novelty re-ranking step, we are able not only to obtain a more visually acceptable set of images, but also to describe with pictures everything that is happening during the whole event. Thus, we are able to avoid focusing on a single set of high relevant images that picture the same concept, activity or background.

²If the informativeness score of a sample is below the threshold, it will be considered non-informative.

Table 1: Best validation accuracy on each set and the respective number of training iterations performed to achieve the results. *SubjX* represents the anonymized sets for a given Subject.

	<i>SubjA</i> ₁	<i>SubjA</i> ₂	<i>SubjB</i> ₁	<i>SubjB</i> ₂	<i>SubjB</i> ₃	<i>SubjC</i> ₁	<i>SubjD</i> ₁	<i>SubjE</i> ₁	<i>SubjF</i> ₁	<i>SubjF</i> ₂
Accuracy	0.759	0.841	0.805	0.795	0.799	0.837	0.805	0.867	0.795	0.897
Iteration #	3,600	1,000	3,600	3,400	3,800	2,600	2,200	3,600	2,400	18,000

**Figure 8: Three examples of the top 5 images obtained before introducing diversity (uneven rows) and after introducing it (even rows).**

5.3 Ranking Summary Quality

The first round of blind taste tests posed two questions for each event. Firstly, experts decided whether each presented summary was representative, and secondly, they chose the preferred one. Given the question *Is the summary representative of the event?*, in 95.74% the experts agreed with our solution, using the ImageNet features to compute the similarity in the novelty-based re-ranking. Using the other alternative, features from Places CNN, the obtained result was 94.33%. We conclude that our approach generates representative summavery high portion of events.

The second question asked was: *Which summary do you prefer?* and experts had to choose their preferred summary, allowing just one answer. The solution based on Imagenet features was chosen in 59.57% of teh cases, while the one based on Places was selected

53.19% of the times. The total adds more than 100% because some summaries were identical for both configurations.

The second round of evaluations aimed to compare our solution based on ImageNet features with a baseline of uniform sampling and and upper-bound defined by the summaries in the ground truth. Experts were asked to *grade each visual summary from 1 (worse) to 5 (best)*, so we could compute the Mean Opinion Score (MOS) [35] of each solution. We adopted MOS as a metric given the highly subjective and complex nature of the task.

Table 2: Mean Opinion Score for ImageNet, ground-truth and uniform sampling summaries.

Our solution	Ground-truth	Uniform Sampling
4,57	4,94	3,99

The performance obtained by uniform sampling (3.99/5) is truly commendable, since this score can be interpreted as a *good* solution. The results obtained with our solution are also satisfactory, because they are closer to the ground-truth than to the uniform sampling. Experts have shown coherence with their ground-truth giving a 4.94 of Mean Opinion Score.

6 CONCLUSIONS

In this paper, we presented a novel approach for semantic summarization of egocentric photo stream events, constructed as a ranked list of general semantic criteria. Our method is based on two main criteria: relevance to optimize semantic diversity in the summary, and novelty to avoid redundancy in the final result. After applying a pre-processing step to filter non-informative images through a new CNN-based method, relevance diversity-aware ranking is obtained by integrating state of the art techniques for saliency detection, object recognition and face detection. This list is re-ranked to reduce redundancy so that each image of the truncated list differs as much as possible from its predecessors. We proposed a new soft metric to rank the informative frames and construct the final summary that does not penalize summaries equivalent (very similar, but not coinciding exactly) to the ground truth. Experimental results indicate high acceptance and satisfaction of psychologists achieving mean opinion score of 4.57 out of 5.0.

ACKNOWLEDGMENTS

This research was supported by contracts SGR1219 and SGR1421 by the Catalan AGAUR office, and TIN2012-38187-C03-01 and TEC2013-43935-R by the Spanish Ministerio de Economía y Competitividad and the European Regional Development Fund (ERDF). We acknowledge the support of NVIDIA Corporation for the donation of GPUs.

REFERENCES

- [1] Omid Aghazadeh, Josephine Sullivan, and Stefan Carlsson. 2011. Novelty detection from an ego-centric perspective. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. IEEE, 3297–3304.
- [2] Michael Blighe, Aiden Doherty, Alan F. Smeaton, and Noel E. O'Connor. 2008. Keyframe Detection in Visual Lifelogs. In *Proceedings of the 1st International Conference on Pervasive Technologies Related to Assistive Environments (PETRA '08)*. ACM, New York, NY, USA, Article 55, 2 pages. <https://doi.org/10.1145/1389586.1389652>
- [3] Marc Bolaños, Ricard Mestre, Estefanía Talavera, Xavier Giró-i Nieto, and Petia Radeva. 2015. Visual Summary of Egocentric Photostreams by Representative Keyframes. *arXiv preprint arXiv:1505.01130* (2015).
- [4] Jaime Carbonell and Jade Goldstein. 1998. The Use of MMR, Diversity-based Reranking for Reordering Documents and Producing Summaries. In *Proceedings of the 21st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '98)*. ACM, New York, NY, USA, 335–336. <https://doi.org/10.1145/290941.291025>
- [5] Charles LA Clarke, Maheedhar Kolla, Gordon V Cormack, Olga Vechtomova, Azin Ashkan, Stefan Büttcher, and Ian MacKinnon. 2008. Novelty and diversity in information retrieval evaluation. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 659–666.
- [6] Duc-Tien Dang-Nguyen, Luca Piras, Giorgio Giacinto, Giulia Boato, and FGB De Natale. 2014. Retrieval of Diverse Images by Pre-filtering and Hierarchical Clustering. *Working Notes of MediaEval* (2014).
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 248–255.
- [8] Thomas Deselaers, Tobias Gass, Philippe Dreuw, and Hermann Ney. 2009. Jointly Optimising Relevance and Diversity in Image Retrieval. In *Proceedings of the ACM International Conference on Image and Video Retrieval (CIVR '09)*. ACM, New York, NY, USA, Article 39, 8 pages. <https://doi.org/10.1145/1646396.1646443>
- [9] Thomas Deselaers, Tobias Gass, Philippe Dreuw, and Hermann Ney. 2009. Jointly optimising relevance and diversity in image retrieval. In *Proceedings of the ACM international conference on image and video retrieval*. ACM, 39.
- [10] Aiden R. Doherty, Daragh Byrne, Alan F. Smeaton, Gareth J.F. Jones, and Mark Hughes. 2008. Investigating Keyframe Selection Methods in the Novel Domain of Passively Captured Visual Lifelogs. In *Proceedings of the 2008 International Conference on Content-based Image and Video Retrieval (CIVR '08)*. ACM, New York, NY, USA, 259–268. <https://doi.org/10.1145/1386352.1386389>
- [11] Joydeep Ghosh, Yong Jae Lee, and Kristen Grauman. 2012. Discovering important people and objects for egocentric video summarization. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1346–1353.
- [12] Boqing Gong, Wei-Lun Chao, Kristen Grauman, and Fei Sha. 2014. Diverse Sequential Subset Selection for Supervised Video Summarization. In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger (Eds.). Curran Associates, Inc., 2069–2077. <http://papers.nips.cc/paper/5413-diverse-sequential-subset-selection-for-supervised-video-summarization.pdf>
- [13] Michael Gygli, Helmut Grabner, Hayko Riemschneider, and Luc Van Gool. 2014. Creating summaries from user videos. In *European conference on computer vision*. Springer, 505–520.
- [14] Steve Hodges, Lyndsay Williams, Emma Berry, Shahram Izadi, James Srinivasan, Alex Butler, Gavin Smyth, Narinder Kapur, and Ken Wood. 2006. SenseCam: A Retrospective Memory Aid. In *Proceedings of the 8th International Conference on Ubiquitous Computing (UbiComp'06)*. Springer-Verlag, Berlin, Heidelberg, 177–193. https://doi.org/10.1007/11853565_11
- [15] Judy Hoffman, Sergio Guadarrama, Eric Tzeng, Jeff Donahue, Ross B. Girshick, Trevor Darrell, and Kate Saenko. 2014. LSDA: Large Scale Detection Through Adaptation. *CoRR abs/1407.5035* (2014). <http://arxiv.org/abs/1407.5035>
- [16] Bogdan Ionescu, Adrian Popescu, Mihai Lupu, Alexandru I. Gnsca, and Henning Müller. 2014. Retrieving diverse social images at mediaeval 2014: Challenge, dataset and evaluation. In *MediaEval 2014 Workshop, Barcelona, Spain*.
- [17] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*. ACM, 675–678.
- [18] Amornched Jinda-Apiraksa, Jana Machajdik, and Robert Sablatnig. 2012. A keyframe selection of lifelog image sequences. *Erasmus Mundus M. Sc. in Visions and Robotics thesis, Vienna University of Technology (TU Wien)* (2012).
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [20] Alex Kulesza and Ben Taskar. 2012. Determinantal point processes for machine learning. *arXiv preprint arXiv:1207.6083* (2012).
- [21] Matthew L. Lee and Anind K. Dey. 2008. Lifelogging Memory Appliance for People with Episodic Memory Impairment. In *Proceedings of the 10th International Conference on Ubiquitous Computing (UbiComp '08)*. ACM, New York, NY, USA, 44–53. <https://doi.org/10.1145/1409635.1409643>
- [22] Zheng Lu and Kristen Grauman. 2013. Story-driven summarization for egocentric video. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*. IEEE, 2714–2721.
- [23] S. Mann. 1998. 'WearCam' (The wearable camera): personal imaging systems for long-term use in wearable tetherless computer-mediated reality and personal photo/video-graphic memory prosthesis. In *Wearable Computers, 1998. Digest of Papers. Second International Symposium on*. 124–131. <https://doi.org/10.1109/ISWC.1998.729538>
- [24] W. W. Mayol, B. J. Tordoff, and D. W. Murray. 2002. Wearable Visual Robots. *Personal Ubiquitous Comput.* 6, 1 (Jan. 2002), 37–48. <https://doi.org/10.1007/s007790200004>
- [25] Mark Montague and Javed A Aslam. 2001. Relevance score normalization for metasearch. In *Proceedings of the tenth international conference on Information and knowledge management*. ACM, 427–433.
- [26] Junting Pan, Kevin McGuinness, and Xavier Giró-i Nieto. 2016. End-to-end Convolutional Network for Saliency Prediction. In *Submitted to CVPR*.
- [27] P. Piasek, K. Irving, and A.F. Smeaton. 2011. SenseCam intervention based on Cognitive Stimulation Therapy framework for early-stage dementia. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2011 5th International Conference on*. 522–525.
- [28] Paulina Piasek, Alan F Smeaton, et al. 2014. Using lifelogging to help construct the identity of people with dementia. (2014).
- [29] Sachan Priyamvada Rajendra and N Keshaveni. 2014. A survey of automatic video summarization techniques. *International Journal of Electronics, Electrical and Computational System 2* (2014).
- [30] M Elena Renda and Umberto Straccia. 2003. Web metasearch: rank vs. score based rank aggregation methods. In *Proceedings of the 2003 ACM symposium on Applied computing*. ACM, 841–846.
- [31] Abigail J. Sellen, Andrew Fogg, Mike Aitken, Steve Hodges, Carsten Rother, and Ken Wood. 2007. Do Life-logging Technologies Support Memory for the Past?: An Experimental Study Using Sensecam. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 81–90. <https://doi.org/10.1145/1240624.1240636>
- [32] Kai Song, Yonghong Tian, Wen Gao, and Tiejun Huang. 2006. Diversifying the Image Retrieval Results. In *Proceedings of the 14th Annual ACM International Conference on Multimedia (MULTIMEDIA '06)*. ACM, New York, NY, USA, 707–710. <https://doi.org/10.1145/1180639.1180789>
- [33] Aimee Spector, Lene Thorgrimsen, Bob Woods, Lindsay Royan, Steve Davies, Margaret Butterworth (deceased), and Martin Orrell. 2003. Efficacy of an evidence-based cognitive stimulation therapy programme for people with dementia. *The British Journal of Psychiatry* 183, 3 (2003), 248–254. <https://doi.org/10.1192/bjp.183.3.248>
- [34] Eleftherios Spyromitros-Xioufis, Symeon Papadopoulos, Yiannis Kompatsiaris, and Ioannis Vlahavas. 2014. SocialSensor: Finding Diverse Images at MediaEval 2014. (2014).
- [35] Robert C Streijl, Stefan Winkler, and David S Hands. 2016. Mean opinion score (MOS) revisited: methods and applications, limitations and alternatives. *Multimedia Systems* 22, 2 (2016), 213–227.
- [36] Estefanía Talavera, Mariella Dimiccoli, Marc Bolaños, Maedeh Aghaei, and Petia Radeva. 2015. R-clustering for egocentric video segmentation. In *Pattern Recognition and Image Analysis*. Springer, 327–336.
- [37] Reinier H. van Leuken, Lluís Garcia, Ximena Olivares, and Roelof van Zwol. 2009. Visual Diversification of Image Search Results. In *Proceedings of the 18th International Conference on World Wide Web (WWW '09)*. ACM, New York, NY, USA, 341–350. <https://doi.org/10.1145/1526709.1526756>
- [38] Cheng Xiang Zhai, William W Cohen, and John Lafferty. 2003. Beyond independent relevance: methods and evaluation metrics for subtopic retrieval. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 10–17.
- [39] Xiangxin Zhu and Deva Ramanan. 2012. Face detection, pose estimation, and landmark localization in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 2879–2886.