University of Groningen

Sentiment Recognition in Egocentric Photostreams

Talavera Martínez, Estefanía; Strisciuglio, Nicola; Petkov, Nicolai; Radeva, Petia

Published in:
Pattern Recognition and Image Analysis

DOI:
[10.1007/978-3-319-58838-4](https://doi.org/10.1007/978-3-319-58838-4)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Early version, also known as pre-print

Publication date:
2017

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Talavera Martínez, E., Strisciuglio, N., Petkov, N., & Radeva, P. (2017). Sentiment Recognition in Egocentric Photostreams. In L. A. Alexandre, J. Salvador Sanchez, & J. M. F. Rodrigues (Eds.), *Pattern Recognition and Image Analysis: Proceedings* (Lecture Notes in Computer Science ; Vol. 10255). Springer. <https://doi.org/10.1007/978-3-319-58838-4>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Sentiment Recognition in Egocentric Photostreams

Estefania Talavera^{1,2}, Nicola Strisciuglio¹, Nicolai Petkov¹, and Petia Radeva^{2,3}

¹ Intelligent Systems Group, University of Groningen, The Netherlands
e.talavera.martinez@rug.nl,

² Department of Mathematics and Computer Science, University of Barcelona, Spain

³ Computer Vision Center, Barcelona, Spain

Abstract. Lifelogging is a process of collecting rich source of information about daily life of people. In this paper, we introduce the problem of sentiment analysis in egocentric events focusing on the moments that compose the images recalling positive, neutral or negative feelings to the observer. We propose a method for the classification of the sentiments in egocentric pictures based on global and semantic image features extracted by Convolutional Neural Networks. We carried out experiments on an egocentric dataset, which we organized in 3 classes on the basis of the sentiment that is recalled to the user (positive, negative or neutral).

Keywords: egocentric photos, lifelogging, sentiment image analysis

1 Introduction

Mental imagery is the process in which the feeling of an experience is imagined by a person in the absence of external stimuli. It has been assumed by therapists to be directly related with emotions [7], opening some questions when images describing past moments of our lives are available: *Can an image make the process of mental imagery easier?* or *Can specific images help us to retrieve or imply feelings and moods?*

Lifelogging is a recent trend consisting in constructing a digital collection from an egocentric point of view of the events of a person that wears a recording device. It is a tool for the analysis of the lifestyle of users, since it provides objective information of what happened during different moments of the day, and a powerful tool for memory enhancement [11]. Using wearable cameras, each day up to 2000 egocentric photostreams are usually recorded, i.e. up to 70000 per month. A lot of these images are redundant, non-informative or routine and thus without special value for the wearer to be preserved. Usually, users are interested in keeping special moments, images with sentiments that will allow them in the future to re-live the personal moments captured by the camera. An automatic tool for sentiment analysis of egocentric images is of high interest to make possible the processing of the big collection of lifelogging data and keeping out just the images of interest i.e. of high charge of positive sentiments.

However, the automatic sentiment image analysis is a complicated task first of all, because of the lack of clear definition of it. There is no consensus between

Fig. 1: Examples of Positive (green), Negative (red) and Neutral (yellow) images.

the different sentiment ontologies in the literature. Table 1 illustrates the ambiguity of the problem, reporting several sentiments ontology in images. The first group [13,21,18] assigns 8 main sentiments as excitement, awe or sadness to the images with assigned discrete positive (1) and negative (-1) sentiment value. The second group [4,10] defines a different set of sentiments as valence or arousal and discrete positive (1), neutral (0) or negative (-1) values assigned to the images according to the sentiments. In contrast, the third group [14] assigns up to 17 sentiments (6 basics and 9 complex) and each image of the dataset is assigned a continuous value in a scale from 1 to 4. Given the ambiguity of the semantic sentiment assignment, with labels difficult to classify into positive or negative sentiments, the last group [1] defines up to 3244 Adjective Noun Pairs (ANP) (e.g. 'beautiful_girl') and assigns a continuous sentiment value in a range of $[-2,2]$ to them. The main idea is that the same object according to its appearance has positive or negative sentiment value like '*angry_dog*' (-1.55) and '*adorable_dog*' (+1.45). A natural question is until which extent the 3244 ANPs represent a scene captured by the image, taking into account the difficulty to detect them automatically (Mean average accuracy $\sim 25\%$).

Given the difficulty of image sentiment determination, ambiguity and lack of consensus in the bibliography, added by the difficulty of the egocentric images, we focus on the image sentiment as a discrete value expressing a ternary sentiment value (positive (1), negative (-1) or neutral (0) value) similar to [20]. Egocentric data is of special difficulty, since we do not observe the wearer and his/her, i.e. from facial or corporal expressions, but rather from the perspective of what the user sees. Moreover, in real life fortunately, negative emotions have much less prevalence than neutral and positive, that makes very difficult to have enough examples of negative egocentric images and events. Thus, the problem we address in this article is what effect an egocentric image or event has on an observer (positive, neutral or negative) (see Fig.1), instead of attempting to specify an explicit semantic image sentiment like sadness; and how to develop automatic tool for sentiment value detection (positive, vs. neutral vs. negative) and egocentric dataset in order to validate its results. Going further, in contrast to the published work, we claim to automatically analyse the sentiment value of egocentric events i.e. a group of sequential images that represents the same scene. In the case of egocentric images, the probability that a single image describes an event is low; there are a lot of images that just capture wall, sky, ground or partially objects. For this reason, we are interested to automatically discover how the event captured by the camera influences the observer, that is to automatically determine the ternary sentiment values of the events, which are

Table 1: Different image sentiment ontologies.

DataSets	Source	Images	Semantic sentiment labels	Sentiment Values
Abstract & Artphoto [13]		280 & 806	positive: contentment, amusement, excitement, awe, negative: sadness, fear, disgust, and anger	{1,-1}
You's Dataset [21]	Flickr Instagr	23000	positive: contentment, amusement, excitement, awe, negative: sadness, fear, disgust, and anger	{1,-1}
CASIA-WebFace [18]		494k	anger, disgust, fear happy, neutral, sad, surprise	[1,0,-1]
IAPS[10]		1182	valence, arousal, and dominance	[1,7]
GAPED [4]		732	valence, arousal, and normative significance	{1,0,-1}
EmoReact [14]	Youtube	1102 clips	17 sentiments: 6 basic emotions (positive: happiness, surprise, negative: sadness, fear, disgust, and anger), and 9 complex emotions: (curiosity, uncertainty, excitement, attentiveness, exploration, confusion, anxiety, embarrassment, frustration).	[1,4]
VSO + TwitterIm[1]	Flickr Twitter	0.5M	Not, but Adjective Noun Pairs (3244)	Flickr[-2,2] Twitter[-1,1]
You_RobustSet [20]	Twitter	1269	Non-semantic labels: Positive and Negative	{1,-1}
UBRUG-EgoSenti*	Wearable Camera	12088	Non-semantic labels: Positive, Neutral and Negative	{1,0,-1}

richer in information and involve the whole moment's experience. For example, an event being in a dark and narrow, grey space would influence negatively, a routine scene like working in the wearer's office could influence the observer neutrally and an event where the wearer has spent some time with friends in a nice outdoor space could influence positively to the observer.

Automatic sentiment analysis from images is a recent research field. In the literature, sentiment recognition in conventional images has been approached by computing and combining visual, textual, or audio features [14,15,17,19]. Other characteristics, such as facial expressions have also been used for sentiment prediction [23]. The combination of visual and textual features extracted from images is possible due to the wide use of online social media and microblogs, where images are posted accompanied by short comments. Therefore, multimodal approaches were proposed, where both sources of information are merged [17,19] for automatic sentiment value detection.

Recently, with the outstanding performance of Convolutional Neural Networks (CNN), several approaches to sentiment analysis relied on deep learning techniques for classification and/or features extraction combined with other networks or methods [2,12,21,22]. The work in [21] applies fine-tuning on the AlexNet to classify the 8 emotions: sadness, angry, content, etc. In contrast, in

[2] they propose to fine-tune CaffeNet with oversampling to classify into Positive or Negative sentiments. In [12] a novel transformation of image intensities to 3D spaces is proposed to reduce the amount of data required to effectively train deep CNN models. In [22] the authors use logistic regression to classify into 3 sentiments using CNN features. In [3], the authors perform a fine-tuning on a CNN model and modify the last layer to classify 2089 ANPs. However, no work has addressed the sentiment image and event analysis in egocentric datasets.

To address the egocentric data sentiment analysis, we propose to combine semantic concepts in terms of ANPs, given that they have sentiment values associated [1], with general visual features extracted from a CNN [9]. ANPs represent a finite subset of concepts present in the image, so they bring strong sentiment value, but can not ensure to cover the whole image content. Visual features extracted by CNNs can help to summarize the whole image content in an intermediate level. We test our method on a new egocentric dataset of 12088 pictures with ternary sentiment values acquired from 3 users and 20 days. A very preliminary stage of this work has been presented in [16].

Therefore, our contributions here are three-fold: a) a model for ternary sentiment value analysis in egocentric images, b) extension of the approach to egocentric events, and c) the first egocentric sentiment value dataset from 12088 images covering 20 days of 3 persons.

The paper is organized as follows. We describe the proposed approach and the dataset in Sections 2 and 3, respectively. In Section 4, we describe the experimental setup, the quantitative and qualitative evaluation, and discuss our findings. Finally, Section 5 draws conclusions and outlines future works.

2 Proposed Method

In this section, we describe the proposed method for sentiment recognition from egocentric photostreams, which is based on visual (extracted by CNN) and semantic (in terms of ANPs) features extracted from the images. An architectural overview of the proposed system is depicted in Fig. 2.

a) Temporal Segmentation: Given that egocentric images have smaller field of view and thus do not capture entirely the context of the event, we need to detect the events of the days. To this aim, we apply the SR-Clustering algorithm for temporal segmentation of photostreams [5]. The clustering procedure is performed on an image representation that combines visual features extracted by a CNN with semantic features in terms of visual concepts extracted by Imagga’s auto-tagging technology (<http://www.imagga.com/solutions/auto-tagging.html>).

b) Features Extraction: For the computation of the semantic features in terms of the ANPs, we use the DeepSentiBank Network [3]. Given an image, the DeepSentiBank network considers the 2089 best performing ANPs. Applying the DeepSentiBank on them gives a 2089-D feature vector, where the feature values correspond to the ANPs likelihood in the image. These values are multiplied by the sentiment value associated to the concepts. Note that each ANP has a positive or negative sentiment value assigned, but not 0 for a neutral sentiment.

However, the 2089 ANPs not necessarily have the power to explain the "richness" of any scene in an image. Hence, we integrate the ANPs feature vector with a feature descriptor provided by the penultimate layer of a CNN [9] that summarizes the whole context of the image. The resulting feature vector is composed by 4096 features. We combine the ANPs and the CNN feature vectors into a 6185-D feature vector, in order to construct a more reliable and rich image representation that relates image semantics expressed by the ANPs with clear sentiment value with the CNN cues as an intermediate image representation. We apply the Signed Root Normalization (SRN) to transform the CNN feature vectors to a more uniformly distributed space followed by a l_2 -normalization [24].

Fig. 2: Architecture of the proposed method. (a) Temporal segmentation of the photostream into events. (b) CNN and ANPs features are extracted from the images and (c) used as input to the trained multi-class SVM model. (d) The model labels the input image as Positive, Neutral or Negative.

c) Classification: We use the proposed feature vectors to train a multi-class SVM classifier due to its high generalization capability [8]. This is ensured by the SVM learning algorithm that finds a separation hyperplane that maximizes the separation margin between the classes. We employ a 1-vs-all design for the multi-class problem, as suggested in [6]. The cardinality of the classes in the proposed dataset is not balanced, which affects the computation of the training error cost. In order to classify an event, we use a majority vote on the image level classification output.

3 Dataset

We collected a dataset of 12471 egocentric pictures, which we call UBRUG-EgoSenti. The users were asked to wear a Narrative Clip Camera, which takes a picture every 30 seconds, hence each day around 1500 images are collected for processing. The images have a resolution of 5MP and JPG format.

We organize the images into events according to the output of the SR-clustering algorithm [5]. From the originally recorded data, we discarded those events that are composed of less than 6 images, so obtaining a dataset composed of 12088 images grouped in a total of 233 events, with an average of 51.87 images per event and std of 52.19. We manually labelled the events following how the user felt while reviewing them by assigning *Positive*, *Negative* and *Neutral* values to them, some examples of which are given in Fig 1. The dataset, for

which the details are in Table 2, is publicly available and can be downloaded from: <http://www.ub.edu/cvub/dataset/>.

Table 2: Description of the UBRUG-EgoSenti dataset.

Class	Images	#Events	Mean Im Event	Std Im Event
Positive	4737	83	57.07	52.34
Neutral	6169	107	57.65	57.18
Negative	1182	43	27.49	26.44
Total	12088	233	51.88	52.19

4 Experiments

4.1 Evaluation and results

We carried out 10-fold cross-validation. Events from different classes are uniformly distributed among the various folds, which are thus independent from each other. We evaluated the performance of the proposed system on single images and at event level. For the UBRUG-Senti dataset, the groundtruth labels are given at event level. All the images that compose a certain event, are considered as having the same label of such event. Given an event composed of M images, we aggregate the M classification decisions by majority vote. We measure the performance results of our method by computing the average accuracy.

Table 3: Performance results achieved at image and event level.

	Image Classification				Event Classification			
	Pos	Neg	Neu	All	Pos	Neg	Neu	All
	mean			mean std	mean			mean std
Semantic Features	59.2	42.4	44.4	48.67 22.87	71.2	42	47.3	53.50 30.77
CNN Features	70	61.3	45.7	59.00 22.80	80.8	71	48.9	66.90 27.67
Semantic+CNN Features	72	60.8	46	59.60 23.17	82.1	73.5	48.9	68.17 30.07

In Table 3, we report the results achieved by the proposed methods at image and event level. We achieved an average image classification rate of 59.60% with a standard deviation of 23.17, when we apply the proposed method. The average event classification rate is 68%, when the proposed features are employed, which corresponds to 82%, 73.5% and 49% for positive, negative and neutral events, respectively. Up to our knowledge, unfortunately, there is no work in the literature on egocentric image sentiment recognition neither event sentiment recognition to compared with. Even the works on image sentiment analysis in conventional images [2,12,21,22] use different datasets and objectives (8 semantic sentiments vs. binary or ternary sentiment values) that make difficult their direct comparison. Fig. 3 shows some example results. As can be seen, the algorithm learns to classify events with presence of routine objects into *neutral* events. Events wrongly classified as *neutral* are shown in Fig. 3(left) and Fig. 3(middle). As an example, the last row of Fig. 3(left) is classified as *neutral*, probably due to the presence of the *pc* in the image, while it was manually labelled as *positive*, because it shows social interactions. As for Fig. 3(left) and Fig. 3(right), events were mislabelled as *negative* probably due to the "homogeneity" and "greyness" of the images within the events, e.g. events were considered as *negative* when most of the information in the image corresponded to the asphalt of the road.

Fig. 3: Examples of the automatic event sentiment classification. The events are grouped based on the sentiment defined by the user: (right) Positive, (middle) Negative, and (left) Neutral. The events frame colour corresponds to the label given by the model: Positive (green), Negative (red) and Neutral (yellow).

4.2 Discussion

Sentiments recognition from an image or a collection of images is a difficult process due to its ambiguity. A challenge in the model construction for sentiment recognition consists in taking into account the bias due to the subjective interpretation of images by different users. Furthermore, the boundaries between neutral/positive and neutral/negative sentiments are not clearly defined. A *neutral* feeling is difficult to interpret. From the results, we observe that *neutral* events are the most challenging ones to classify. Another challenging aspect concerns the grouping of image sentiments into event sentiment, since events can have non-uniform sentiments.

A further step towards better understanding of the image and sentiment analysis is needed, due to the subjectivity of what an image can recall to different persons. To this aim, having annotations by different persons is critical to evaluate the inter- and intra-observer variability.

From the results, the intuition that we get is that non-routine events and specially when moments are social, have a higher probability of being positive. In contrast, routine events will most probably be considered as neutral. Negative events as accidents have low prevalence to be learned. Yet, hostile and empty environments could lead to negative sentiments too. Future works will address the study of emotional events and their relation to daily routine.

5 Conclusions

In this work, we propose, for the first time, a system and a dataset for egocentric sentiment image and event recognition based on the extraction of CNN and semantic features with sentiment value associated. We introduced a new labelled dataset of egocentric images composed of 233 events, grouping 12088 images, from 20 days of 3 users grouped. We presented preliminary results, obtaining an average events and image sentiment accuracy of 68.17% and 58.60%, with std of 30.07% and 23.17%, respectively.

Acknowledgements

This work was partially founded by TIN2015-66951-C2, SGR 1219, CERCA, *ICREA Academia'14* and Grant 20141510 (Marat6TV3). The funders had no role in the study design, data collection, analysis, and preparation of the manuscript.

References

1. D. Borth, R. Ji, T. Chen, T. Breuel, and S.-F. Chang. Large-scale visual sentiment ontology and detectors using adjective noun pairs. *ACM*, pages 223–232, 2013.
2. V. Campos and et al. Diving Deep into Sentiment: Understanding Fine-tuned CNNs for Visual Sentiment Prediction. *ASM*, pages 57–62, 2015.
3. T. Chen, D. Borth, T. Darrell, and S.-F. Chang. DeepSentiBank: Visual Sentiment Concept Classification with Deep Convolutional Neural Networks. page 7, 2014.
4. E. S. Dan-Glauser and K. R. Scherer. The Geneva affective picture database (GAPED): a new 730-picture database focusing on valence and normative significance. *Behavior research methods*, 43(2):468–77, 2011.
5. M. Dimiccoli, E. Talavera, S. G. Nikolov, and P. Radeva. SR-Clustering: Semantic Regularized Clustering for Egocentric Photo Streams Segmentation. 2015.
6. P. Foggia, N. Petkov, A. Saggese, N. Strisciuglio, and M. Vento. Reliable detection of audio events in highly noisy environments. *PRL*, 65(1):22–28, 2015.
7. E. A. Holmes and et al. Positive Interpretation Training: Effects of Mental Imagery Versus Verbal Training on Positive Mood. *Behavior Therapy*, 37(3):237–247, 2006.
8. T. Joachims. Estimating the Generalization Performance of a SVM efficiently. *ICML*, pages 431–438, 2000.
9. A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *NIPS*, pages 1–9, 2012.
10. P. Lang, M. Bradley, and B. Cuthbert. International Affective Picture System (IAPS): Technical Manual and Affective Ratings. *NIMH*, pages 39–58, 1997.
11. M. L. Lee and A. K. Dey. Lifelogging memory appliance for people with episodic memory impairment. *UbiComp*, 2008.
12. G. Levi and T. Hassner. Emotion Recognition in the Wild via Convolutional Neural Networks and Mapped Binary Patterns. *ICMI*, pages 503–510, 2015.
13. J. Machajdik and A. Hanbury. Affective Image Classification using Features inspired by Psychology and Art Theory. *ICM*, pages 83–92, 2010.
14. B. Nojavanasghar and et al. EmoReact: A Multimodal Approach and Dataset for Recognizing Emotional Responses in Children. *ICMI 2016*, pages 137–144, 2016.
15. S. Poria and et al. Fusing audio, visual and textual clues for sentiment analysis from multimodal content. *Neurocomputing*, 174:50–59, 2014.
16. E. Talavera, P. Radeva, and N. Petkov. Towards Egocentric Sentiment Analysis. In *16th International Conference on Computer Aided Systems Theory*, 2017.
17. M. Wang, D. Cao, L. Li, S. Li, and R. Ji. Microblog Sentiment Analysis Based on Cross-media Bag-of-words Model. *ICIMCS*, pages 76–80, 2014.
18. D. Yi, Z. Lei, S. Liao, and S. Z. Li. Learning Face Representation from Scratch. *arXiv*, 2014.
19. Q. You and A. Et. Cross-modality Consistent Regression for Joint Visual-Textual Sentiment Analysis of Social Multimedia. *Wsdm*, pages 13–22, 2016.
20. Q. You and et al. Robust Image Sentiment Analysis using Progressively Trained and Domain Transferred Deep Networks. *AAAI*, pages 381–388, 2015.
21. Q. You, J. Luo, H. Jin, and J. Yang. Building a Large Scale Dataset for Image Emotion Recognition: The Fine Print and The Benchmark. *CoRR*, 2016.
22. Y. Yu, H. Lin, J. Meng, and Z. Zhao. Visual and Textual Sentiment Analysis of a Microblog Using Deep Convolutional Neural Networks. *Algorithms*, 9(2):41, 2016.
23. J. Yuan and et al. Sentitribute : Image Sentiment Analysis from a Mid-level Perspective Categories and Subject Descriptors. *WISDOM*, pages 101–108, 2013.
24. L. Zheng, S. Wang, F. He, and Q. Tian. Seeing the Big Picture: Deep Embedding with Contextual Evidences. page 10, 2014.