Multimodal sentiment analysis : addressing key issues and setting up the baselines

Poria, Soujanya; Majumder, Navonil; Hazarika, Devamanyu; Cambria, Erik; Gelbukh, Alexander; Hussain, Amir

2018

Poria, S., Majumder, N., Hazarika, D., Cambria, E., Gelbukh, A., & Hussain, A. (2018). Multimodal sentiment analysis : addressing key issues and setting up the baselines. IEEE Intelligent Systems, 33(6), 17-25. doi:10.1109/MIS.2018.2882362

<https://hdl.handle.net/10356/143239>

<https://doi.org/10.1109/MIS.2018.2882362>

© 2018 IEEE (published by the IEEE Computer Society). Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: <https://doi.org/10.1109/MIS.2018.2882362>.

Downloaded on 26 Apr 2023 07:48:37 SGT

Multimodal Sentiment Analysis: Addressing Key Issues and Setting up the Baselines

Soujanya Poria, Navonil Majumder, Devamanyu Hazarika,
Erik Cambria, Alexander Gelbukh and Amir Hussain

Abstract—We compile baselines, along with dataset split, for multimodal sentiment analysis. In this paper, we explore three different deep-learning based architectures for multimodal sentiment classification, each improving upon the previous. Further, we evaluate these architectures with multiple datasets with fixed train/test partition. We also discuss some major issues, frequently ignored in multimodal sentiment analysis research, e.g., role of speaker-exclusive models, importance of different modalities, and generalizability. This framework illustrates the different facets of analysis to be considered while performing multimodal sentiment analysis and, hence, serves as a new benchmark for future research in this emerging field.

I. INTRODUCTION

Emotion recognition and sentiment analysis is opening up numerous opportunities pertaining social media in terms of understanding users preferences, habits, and their contents [11]. With the advancement of communication technology, abundance of mobile devices, and the rapid rise of social media, a large amount of data is being uploaded as video, rather than text [2]. For example, consumers tend to record their opinions on products using a webcam and upload them on social media platforms, such as YouTube and Facebook, to inform the subscribers of their views. Such videos often contain comparisons of products from competing brands, pros and cons of product specifications, and other information that can aid prospective buyers to make informed decisions.

The primary advantage of analyzing videos over mere text analysis, for detecting emotions and sentiment, is the surplus of behavioral cues. Videos provide multimodal data in terms of vocal and visual modalities. The vocal modulations and facial expressions in the visual data, along with text data, provide important cues to better

identify true affective states of the opinion holder. Thus, a combination of text and video data helps to create a better emotion and sentiment analysis model.

Recently, a number of approaches to multimodal sentiment analysis producing interesting results have been proposed [12], [14]. However, there are major issues that remain mostly unaddressed in this field, such as the consideration of context in classification, effect of speaker-inclusive and speaker-exclusive scenario, the impact of each modality across datasets, and generalization ability of a multimodal sentiment classifier. Not tackling these issues has presented difficulties in effective comparison of different multimodal sentiment analysis methods. In this paper, we outline some methods that address these issues and setup a baseline based on state-of-the-art methods. We use a deep convolutional neural network (CNN) to extract features from visual and text modalities.

This paper is organized as follows: Section II provides a brief literature review on multimodal sentiment analysis; Section III briefly discusses the baseline methods; experimental results and discussion are given in Section IV; finally, Section V concludes the paper.

II. RELATED WORK

In 1970, Ekman et al. [6] carried out extensive studies on facial expressions. Their research showed that universal facial expressions are able to provide sufficient clues to detect emotions. Recent studies on speech-based emotion analysis [4] have focused on identifying relevant acoustic features, such as fundamental frequency (pitch), intensity of utterance, bandwidth, and duration.

As to fusing audio and visual modalities for emotion recognition, two of the early works were done by De Silva et al. [5] and Chen et al. [3]. Both works showed that a bimodal system yielded a higher accuracy than any unimodal system.

While there are many research papers on audio-visual fusion for emotion recognition, only a few research works have been devoted to multimodal emotion or sentiment analysis using text clues along with visual and audio modalities. Wollmer et al. [15] fused information

N. Majumder and A. Gelbukh are with the CIC, Instituto Politécnico Nacional, Mexico City, Mexico.

S. Poria and E. Cambria is with the SCSE, Nanyang Technological Institute, Singapore.

D. Hazarika is with the SOC, National University of Singapore, Singapore.

A. Hussain is with the Edinburgh Napier University, UK.

from audio, visual and text modalities to extract emotion and sentiment. Metallinou et al. [9] fused audio and text modalities for emotion recognition. Both approaches relied on feature-level fusion.

In this paper, we study the behavior of the method proposed in [13] in the aspects rarely addressed by other authors, such as speaker independence, generalizability of the models and performance of individual modalities.

III. UNIMODAL FEATURE EXTRACTION

For the unimodal feature extraction, we follow the procedures by bc-LSTM [13].

A. Textual Feature Extraction

We employ convolutional neural networks (CNN) for textual feature extraction. Following [8], we obtain n-gram features from each utterance using three distinct convolution filters of sizes 3, 4, and 5 respectively, each having 50 feature-maps. Outputs are then subjected to max-pooling followed by rectified linear unit (ReLU) activation. These activations are concatenated and fed to a 100 dimensional dense layer, which is regarded as the textual utterance representation. This network is trained at utterance level with the emotion labels.

B. Audio and Visual Feature Extraction

Identical to [13], we use 3D-CNN and openSMILE [7] for visual and acoustic feature extraction, respectively.

C. Fusion

In order to fuse the information extracted from different modalities, we concatenated the feature vectors representative of the given modalities and sent the combined vector to a classifier for the classification. This scheme of fusion is called feature-level fusion. Since, the fusion involved concatenation and no overlapping, merge, or combination, scaling and normalization of the features were avoided. We discuss the results of this fusion in Section IV.

D. Baseline Method

1) *bc-LSTM*: We follow the method bc-LSTM [13] where they used a bi-directional LSTM to capture the context from the surrounding utterances to generate context-aware utterance representation.

2) *SVM*: After extracting the features, we merged and sent to a SVM with RBF kernel for the final classification.

IV. EXPERIMENTS AND OBSERVATIONS

In this section, we discuss the datasets and the experimental settings. Also, we analyze the results yielded by the aforementioned methods.

A. Datasets

1) *Multimodal Sentiment Analysis Datasets*: For our experiments, we used the MOUD dataset, developed by Perez-Rosas et al. [10]. They collected 80 product review and recommendation videos from YouTube. Each video was segmented into its utterances (498 in total) and each of these was categorized by a sentiment label (positive, negative and neutral). On average, each video has 6 utterances and each utterance is 5 seconds long. In our experiment, we did not consider neutral labels, which led to the final dataset consisting of 448 utterances. We dropped the neutral label to maintain consistency with previous work. In a similar fashion, Zadeh et al. [16] constructed a multimodal sentiment analysis dataset called multimodal opinion-level sentiment intensity (MOSI), which is bigger than MOUD, consisting of 2199 opinionated utterances, 93 videos by 89 speakers. The videos address a large array of topics, such as movies, books, and products. In the experiment to address the generalizability issues, we trained a model on MOSI and tested on MOUD. Table I shows the split of train/test of these datasets.

2) *Multimodal Emotion Recognition Dataset*: The IEMOCAP database [1] was collected for the purpose of studying multimodal expressive dyadic interactions. This dataset contains 12 hours of video data split into 5 minutes of dyadic interaction between professional male and female actors. Each interaction session was split into spoken utterances. At least 3 annotators assigned to each utterance one emotion category: *happy*, *sad*, *neutral*, *angry*, *surprised*, *excited*, *frustration*, *disgust*, *fear* and *other*. In this work, we considered only the utterances with majority agreement (i.e., at least two out of three annotators labeled the same emotion) in the emotion classes of *angry*, *happy*, *sad*, and *neutral*. Table I shows the split of train/test of this dataset.

B. Speaker-Exclusive Experiment

Most of the research on multimodal sentiment analysis is performed with datasets having common speaker(s) between train and test splits. However, given this overlap, results do not scale to true generalization. In real-world applications, the model should be robust to speaker variance. Thus, we performed speaker-exclusive experiments to emulate unseen conditions. This time, our train/test splits of the datasets were completely disjoint

Dataset	Train		Test	
	<i>utterance</i>	<i>video</i>	<i>utterance</i>	<i>video</i>
IEMOCAP	4290	120	1208	31
MOSI	1447	62	752	31
MOUD	322	59	115	20
MOSI $\rightarrow$ MOUD	2199	93	437	79

TABLE I: Person-Independent Train/Test split details of each dataset ($\approx 70/30$ % split). Note: $X \rightarrow Y$ represents train: X and test: Y ; Validation sets are extracted from the shuffled train sets using 80/20 % train/val ratio.

with respect to speakers. While testing, our models had to classify emotions and sentiments from utterances by speakers they have never seen before. Below, we elaborate this speaker-exclusive experiment:

- **IEMOCAP:** As this dataset contains 10 speakers, we performed a 10-fold speaker-exclusive test, where in each round exactly one of the speakers was included in the test set and missing from train set. The same SVM model was used as before and accuracy was used as performance metric.
- **MOUD:** This dataset contains videos of about 80 people reviewing various products in Spanish. Each utterance in the video has been labeled as *positive*, *negative*, or *neutral*. In our experiments, we consider only samples with *positive* and *negative* sentiment labels. The speakers were partitioned into 5 groups and a 5-fold person-exclusive experiment was performed, where in every fold one out of the five group was in the test set. Finally, we took average of the accuracy to summarize the results (Table II).
- **MOSI:** MOSI dataset is rich in sentimental expressions, where 93 people review various products in English. The videos are segmented into clips, where each clip is assigned a sentiment score between -3 to $+3$ by five annotators. We took the average of these labels as the sentiment polarity and naturally considered two classes (*positive* and *negative*). Like MOUD, speakers were divided into five groups and a 5-fold person-exclusive experiment was run. For each fold, on average 75 people were in the training set and the remaining in the test set. The training set was further partitioned and shuffled into 80%–20% split to generate train and validation sets for parameter tuning.

1) *Speaker-Inclusive vs. Speaker-Exclusive:* In comparison with the speaker-inclusive experiment, the speaker-exclusive setting yielded inferior results. This is caused by the absence of knowledge about the speakers

during the testing phase. Table II shows the performance obtained in the speaker-inclusive experiment. It can be seen that audio modality consistently performs better than visual modality in both MOSI and IEMOCAP datasets. The text modality plays the most important role in both emotion recognition and sentiment analysis. The fusion of the modalities shows more impact for emotion recognition than for sentiment analysis. Root mean square error (RMSE) and TP-rate of the experiments using different modalities on IEMOCAP and MOSI datasets are shown in Fig. 1.

Fig. 1: Experiments on IEMOCAP and MOSI datasets. The top-left figure shows the RMSE of the models on IEMOCAP and MOSI. The top-right figure shows the dataset distribution. Bottom-left and bottom-right figures present TP-rate on of the models on IEMOCAP and MOSI dataset, respectively.

C. Contributions of the Modalities

As expected, bimodal, and trimodal models have performed better than unimodal models in all experiments. Overall, audio modality has performed better than visual on all datasets. Except for MOUD dataset, the unimodal performance of text modality is substantially better than other two modalities (Fig. 2).

D. Generalizability of the Models

To test the generalization ability of the models, we trained the framework on MOSI dataset in speaker-exclusive fashion and tested with MOUD dataset. From Table III, we can see that the trained model with MOSI dataset performed poorly with MOUD dataset.

This is mainly due to the fact that reviews in MOUD dataset had been recorded in Spanish, so both audio and text modalities miserably fail in recognition, as MOSI dataset contains reviews in English. A more comprehensive study would be to perform generalizability tests

Modality Combination	IEMOCAP		MOUD		MOSI	
	Sp-In	Sp-Ex	Sp-In	Sp-Ex	Sp-In	Sp-Ex
A	66.20	51.52	–	53.70	64.00	57.14
V	60.30	41.79	–	47.68	62.11	58.46
T	67.90	65.13	–	48.40	78.00	75.16
T + A	78.20	70.79	–	57.10	76.60	75.72
T + V	76.30	68.55	–	49.22	78.80	75.06
A + V	73.90	52.15	–	62.88	66.65	62.4
T + A + V	81.70	71.59	–	67.90	78.80	76.66

TABLE II: Accuracy reported for speaker-exclusive (Sp-Ex) and speaker-inclusive (Sp-In) split for Concatenation-Based Fusion. *IEMOCAP*: 10-fold speaker-exclusive average. *MOUD*: 5-fold speaker-exclusive average. *MOSI*: 5-fold speaker-exclusive average. *Legend*: A stands for Audio, V for Video, T for Text.

Fig. 2: Performance of the modalities on the datasets. Red line indicates the median of the accuracy.

Modality Combination	Accuracy	
	SVM	bc-LSTM
T	46.5%	46.9%
V	43.3%	49.6%
A	42.9%	47.2%
T + A	50.4%	51.3%
T + V	49.8%	49.8%
A + V	46.0%	49.6%
T + A + V	51.1%	52.7%

TABLE III: **Cross-dataset results:** Model (with previous configurations) trained on MOSI dataset and tested on MOUD dataset.

on datasets of the same language. However, we were unable to do this for the lack of benchmark datasets. Also, similar experiments of cross-dataset generalization was not performed on emotion detection, given the availability of only a single dataset (IEMOCAP).

E. Comparison among the Baseline Methods

Table IV consolidates and compares performance of all the baseline methods for all the datasets. We evaluated

SVM and bc-LSTM fusion with MOSI, MOUD, and IEMOCAP dataset.

From Table IV, it is clear that bc-LSTM performs better than SVM across all the experiments. So, it is very apparent that consideration of context in the classification process has substantially boosted the performance.

F. Visualization of the Datasets

MOSI visualizations present information regarding dataset distribution within single and multiple modalities (Fig. 3). For the textual and audio modalities, comprehensive clustering can be seen with substantial overlap. However, this problem is reduced in the video and all modalities scenario with structured declustering but overlap is reduced only in multimodal. This forms an intuitive explanation of the improved performance in the multimodal scenario. IEMOCAP visualizations provide insight for the 4-class distribution for uni and multimodal scenario, where clearly the multimodal distribution has the least overlap (increase in red and blue visuals, apart from the rest) with sparse distribution aiding the classification process.

V. CONCLUSION

We have presented useful baselines for multimodal sentiment analysis and multimodal emotion recognition. We also discussed some major aspects of multimodal sentiment analysis problem, such as the performance in the unknown-speaker setting and the cross-dataset performance of the models.

Our future work will focus on extracting semantics from the visual features, relatedness of the cross-modal features and their fusion. We will also include contextual dependency learning in our model to overcome the limitations mentioned in the previous section.

Modality Combination	IEMOCAP		MOUD		MOSI	
	SVM	bc-LSTM	SVM	bc-LSTM	SVM	bc-LSTM
A	52.9	57.1	51.5	59.9	58.5	60.3
V	47.0	53.2	46.3	48.5	53.1	55.8
T	65.5	73.6	49.5	52.1	75.5	78.1
T + A	70.1	75.4	53.1	60.4	75.8	80.2
T + V	68.5	75.6	50.2	52.2	76.7	79.3
A + V	67.6	68.9	62.8	65.3	58.6	62.1
T + A + V	72.5	76.1	66.1	68.1	77.9	80.3

TABLE IV: Accuracy reported for speaker-exclusive classification. *IEMOCAP*: 10-fold speaker-exclusive average. *MOUD*: 5-fold speaker-exclusive average. *MOSI*: 5-fold speaker-exclusive average. *Legend*: A represents Audio, V represents Video, T represents Text.

Fig. 3: T-SNE 2D visualization of MOSI and IEMOCAP datasets when unimodal features and multimodal features are used.

REFERENCES

- [1] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335–359, 2008.
- [2] E. Cambria, H. Wang, and B. White. Guest editorial: Big social data analysis. *Knowledge-Based Systems*, 69:1–2, 2014.
- [3] L. S. Chen, T. S. Huang, T. Miyasato, and R. Nakatsu. Multi-modal human emotion/expressions recognition. In *Proceedings of the Third IEEE International Conference on Automatic Face and Gesture Recognition*, pages 366–371. IEEE, 1998.
- [4] D. Datcu and L. Rothkrantz. Semantic audio-visual data fusion for automatic emotion recognition. *Euromedia*, 2008.
- [5] L. C. De Silva, T. Miyasato, and R. Nakatsu. Facial emotion recognition using multi-modal information. In *Proceedings of ICICS*, volume 1, pages 397–401. IEEE, 1997.
- [6] P. Ekman. Universal facial expressions of emotion. *Culture and Personality: Contemporary Readings/Chicago*, 1974.
- [7] F. Eyben, M. Wöllmer, and B. Schuller. Opensmile: the munich versatile and fast open-source audio feature extractor. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 1459–1462. ACM, 2010.
- [8] Y. Kim. Convolutional neural networks for sentence classification. *CoRR*, abs/1408.5882, 2014.
- [9] A. Metallinou, S. Lee, and S. Narayanan. Audio-visual emotion recognition using gaussian mixture models for face and voice. In *Tenth IEEE International Symposium on ISM 2008*, pages 250–257. IEEE, 2008.
- [10] V. Pérez-Rosas, R. Mihalcea, and L.-P. Morency. Utterance-level multimodal sentiment analysis. In *ACL*, pages 973–982, 2013.
- [11] S. Poria, E. Cambria, R. Bajpai, and A. Hussain. A review of affective computing: From unimodal analysis to multimodal fusion. *Information Fusion*, 37:98–125, 2017.
- [12] S. Poria, E. Cambria, and A. Gelbukh. Deep convolutional neural network textual features and multiple kernel learning for utterance-level multimodal sentiment analysis. In *Proceedings of EMNLP*, pages 2539–2544, 2015.
- [13] S. Poria, E. Cambria, D. Hazarika, N. Majumder, A. Zadeh, and L.-P. Morency. Context-dependent sentiment analysis in user-generated videos. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 873–883, Vancouver, Canada, July 2017. Association for Computational Linguistics.
- [14] S. Poria, I. Chaturvedi, E. Cambria, and A. Hussain. Convolutional MKL based multimodal emotion recognition and sentiment analysis. In *ICDM*, pages 439–448, Barcelona, 2016.
- [15] M. Wollmer, F. Weninger, T. Knaup, B. Schuller, C. Sun, K. Sagae, and L.-P. Morency. Youtube movie reviews: Sentiment analysis in an audio-visual context. *IEEE Intelligent Systems*, 28(3):46–53, 2013.
- [16] A. Zadeh, R. Zellers, E. Pincus, and L.-P. Morency. Multimodal

sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6):82–88, 2016.