



Trend analysis in online data with unsupervised classification and appraisal categories

Valentina Dragos, Jérôme Besombes, Aurélien Mascaro

► To cite this version:

Valentina Dragos, Jérôme Besombes, Aurélien Mascaro. Trend analysis in online data with unsupervised classification and appraisal categories. 2020 IEEE 23rd International Conference on Information Fusion (FUSION), Jul 2020, virtuel, South Africa. pp.1-8, 10.23919/FUSION45008.2020.9190451 . hal-03212292

HAL Id: hal-03212292

<https://hal.science/hal-03212292>

Submitted on 29 Apr 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Trend analysis in online data with unsupervised classification and appraisal categories

Valentina Dragos

ONERA - The French Aerospace Lab

Palaiseau, France

valentina.dragos@onera.fr

Jérôme Besombes

ONERA - The French Aerospace Lab

Palaiseau, France

jerome.besombes@onera.fr

Aurélien Mascaro

ONERA - The French Aerospace Lab

Palaiseau, France

aurelien.mascaro@onera.fr

Abstract—This paper presents a hybrid approach for identifying trends in social media datasets. This approach uses jointly unsupervised methods for text classification and an ontology of appraisal categories. First, data gleaned on social media are classified with unsupervised methods in order to produce more or less homogeneous clusters. Then, topics are detected within each cluster by using Latent Dirichlet Allocation (LDA). At cluster level, appraisal categories are identified thanks to an ontology build according to the principles of the Appraisal Theory. The joint identification of topics and appraisals offers a basis to analyze trending topics in social data. The paper investigates a novel means to detect trending topics in social data by utilizing unsupervised classification methods and focusing on subjective states such as affect, attitude, denial, disapproval, rejection, endorsement or support, associated with each class. Negative or positive polarity and intensity degrees are also identified, thanks to the appraisal ontology. The approach identifies the most dominant trends in social data as associations of topics and appraisal categories. The paper also discusses experiments carried out to detect trends on Twitter collections and the evaluation of their results in the light of manually validated data.

Index Terms—Social media, trend analysis, hybrid AI, opinion and sentiment detection

I. INTRODUCTION

Social networks are data production engines and generate content on a variety of domains [1], from advertising to politics and societal aspects. Their analytics can produce useful insights for those domains. Discussions on social media are also indicators for various phenomena, from opinion shifts to emergency situations, social unrest or disease outbreaks.

This paper addresses the exploration of the cyber-social space as defined by Shets and colleagues in [2], as a network of humans and autonomous agents and their links creating human, autonomous or hybrid communities. For the cyber-physical space, recent examples on disinformation [3] or social manipulation [4] show that phenomena in the cyberspace have the capacity to polarize social views, making social groups form and fracture in online spaces with concrete consequences on real political and social environments. It is then important to have means to effectively mine the cyberspace for emotion, opinion or sentiment detection and effectively analyze the dynamics of such components inside the network and especially the emotional contagion [5] and online propaganda [6].

The analysis of sentiments and opinions is a field of Natural Language Processing (NLP) domain, with the aim of detecting

[7] valence, opinions, emotions and other subjective states from text.

For social data analysis, Twitter become an increasingly popular source of data, with mainly machine and deep learning models being developed for this tasks [8]. Most of the research studies keep the distinction between factual and subjective aspects and focus either on topics [9], [10] and narrative detection [11], [12] or subjectivity specific tasks: detecting polarities [13], sentiments [14] or opinions [15]. Both tasks are difficult given that content and dynamics in social data varies, sometimes in the same conversation. However, meaningful clues hidden in online data are often a combination of topics and subjective aspects and their identification involves analysis of emotions conveyed towards specific topics, how emotions or opinions of groups change over time [16] as well as detecting the mapping between the emotional categories and linguistic instances [17].

The approach developed in this paper tackles the identification of trends in social data with a hybrid approach. First, unsupervised classification is used to detect topics within the collection and to structure the initial data set accordingly; then, an ontology modelling appraisal categories is used to automatically identify appraisal annotations in the classified data. More specifically, topics encapsulated in each individual tweet are detected using the DBSCAN [18] and Latent Dirichlet Allocation (LDA) [19] and the use of ontology at cluster level allows to pool tweets into trending groups and show which appraisal concepts are the most dominant in topics. Experiments were conducted with several data sets crawled directly from online networks and the results were analyzed against a set of manually annotated data.

The rest of this paper is organized as follows: section II discusses related work; section III describes the methods that are used for mining Twitter data sets : (1) unsupervised clustering (2) detection of topics with LDA and (3) detection of opinions with the appraisal ontology. Section IV presents the experiments and results. Limitations of the approach are discussed in section V. Section VI concludes and sketches directions for future work.

II. RELATED WORK

Social data analysis has been actively studied in the past few years, in an effort to both investigate and understand online

content. Solutions were developed for sentiment analysis and opinion detection, although those terms are not accurately defined and are sometimes interchangeable. More recent applications are in the field of security, with emphasis on extremist content detection [20] and understanding propaganda and ideas propagation in the cyberspace [21].

The goal of early approaches was to detect polarity in products [22] or service reviews [23], and the analysis was primarily performed at sentence [24] or document level [25]. The approaches developed range from lexicon and ontology-based methods [26] to supervised machine learning and deep learning techniques that were developed in recent years [27].

Classification algorithms were widely used for sentiment analysis [28], and those techniques mostly depend on feature engineering and manually defined rules and resources, such as dependency and causality relations, n-grams or sentiment lexicons. They leverage the bag-of-words representation to convert the corpus into a term-document matrix, following several pre-processing techniques, such as normalization and stemming.

More recently, approaches based on neural network techniques without feature engineering became popular for social data analysis [24]. Those solutions rely upon embedded structures, such as low dimensional word vectors which contain shallow semantic information. Similar approaches are developed by adopting supervised sequence labelling. Thus, Hidden Markov models and conditional random fields are used by Chen and colleagues to extract aspect and polarity from social data [29]. To alleviate the need to large amount of labelled data for training purposes, unsupervised methods based on topics models [30] or ontologies [26] were also developed.

Although supervised solutions achieve reasonable accuracy for sentiment analysis, difficulties in processing a mix of multi-domain data and the use of manually tagged inputs have been plaguing the robustness of those approaches [31].

From a different perspective, a variety of semantic resources [32] have been designed to capture and model concepts describing subjective engagement, polarity and other attributes specific to opinion detection and sentiment analysis. Based on those resources, techniques were developed which are able to detect instances of their concepts within textual content and to emphasize on words that are relevant to aspects, emotions or opinions [33]. Many of those resources come with limitations as they are designed to achieve rather broad coverage and fail to capture domain-specific standpoints [34].

However, most work in this area focuses on overall opinion detection or sentiment analysis, regardless of the entities, targets or topics mentioned in the content. The algorithms detect sentiments or opinion by assuming a known target. Taking a step further, recent studies by Schoene and de Mel [35] investigate the correlation of topics and emotions while Vijayaraghavan and colleagues addressed the classification of topics along with the sentiment [36]. Following a similar research line, Schmitt et al. employed an end-to-end trainable neural network to detect aspects and classify polarities jointly. Ensuring important volumes, high dynamic and a flexible

API, Twitter is nowadays a popular source of social data and benchmarks are created in order to assess and compare approaches [37].

This paper addresses the detection of trends from social data by investigating the joint identification of topics and opinions within several collections of tweets. The rationale behind this distinct analysis is that opinions expressed towards topics are good indicators of the support behind ideologies and ideas.

Original contributions stem from the combination of unsupervised classification techniques for topics detection and the use of ontology-driven techniques to detect fine-grained appraisal categories expressing support, deny, rejection or endorsement.

III. A HYBRID ARCHITECTURE FOR TREND ANALYSIS

In order to detect trends and to identify the emerging topics associated with those trends we adopt the following approach:

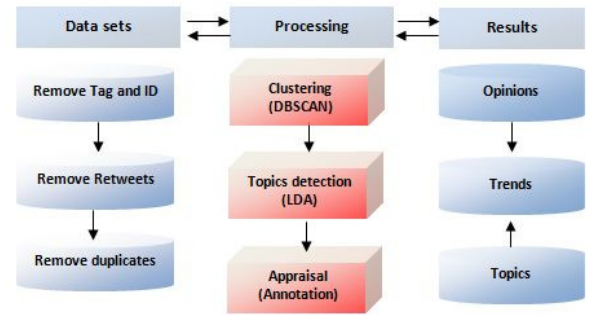


Fig. 1. Hybrid architecture for trend analysis

- Several collections of tweets are crawled and classified thanks to unsupervised techniques;
- Topics are detected within each cluster as to pool tweets into more homogeneous groups and show which topics are the most dominant within each cluster;
- An ontology of appraisal categories is then used to detect appraisal categories in tweets;
- A trend-detection strategy is adopted based on appraisal categories to infer categories detected in topics and to aggregate trends at topic level from appraisals identified at tweet level.

A. Unsupervised learning and feature engineering

The method adopted for data clustering is DBSCAN (Density-Based Spatial Clustering of Applications with Noise) [38]. This technique groups data located in the region with high density of the data space to belong to the same cluster and it discovers clusters with arbitrary shape. The technique also increases cluster homogeneity by filtering out noise.

The DBSCAN algorithm basically requires 2 parameters:

- *eps*: specifies how close points should be to each other to be considered a part of a cluster. If the distance between two points is lower or equal to *eps* value, those points are considered neighbors.

- minPoints: the minimum number of points to form a dense region. For example, if is the minPoints parameter is set as 5, then at least 5 points are needed to form a dense region.

The main advantage of DBSCAN is that the method is not restricted to a set number of clusters during initialization, although the algorithm will determine a number of clusters based on the density of a region.

For this work, DBSCAN has been applied iteratively on disjoint dataset portions and all the original dataset is clustered at the first level. Then, tweets labelled as outliers in the previous level are re-clustered at each subsequent level. In an attempt to discover representative clusters within datasets, the methods avoid clusters containing few tweets. The number of tweets labelled as outliers and thus unclustered is also limited, in order to consider all different posted items. Through addressing these issues, DBSCAN parameters were properly set at each level by using the Elbow heuristic [39].

Several research surrounding Density-based algorithms in Twitter mining show that they are suitable for clustering unstructured data [40]. However, unsupervised classification puts forth a particular challenge because the algorithm groups data into similar categories, without requiring prior understanding of the groups content.

In the context of this work clusters can be of different topics and Latent Dirichlet Allocation(LDA) [41] was used to extract the topics from the vocabulary of previously defined clusters.

LDA is a probabilistic model developed to describe collections of discrete data such as textual corpora composed of words. From a technical standpoint, LDA is a three-level hierarchical Bayesian model, in which each item of a collection is modeled as a finite mixture over an underlying set of topics. Iteratively, each topic is also modeled as an infinite mixture over an underlying set of topic probabilities. In the context of text modeling, the topic probabilities association provides an explicit representation of a document.

Although many clustering models restrict a document to being associated with a single topic, under LDA documents can be associated with multiple topics, as the topic node is sampled repeatedly within the document. The method is also unsupervised and topics are detected without needing to define tags or train data beforehand.

B. The appraisal ontology and subjectivity analysis

Trend detection is based on the cognitive foundations of the appraisal theory and uses the ontology of appraisals to annotate tweets.

The Appraisal Theory [42] is a cognitive frame claiming that people's emotions are elicited by their personal and continuous interpretations, evaluations or appraisals of objects, events and situations. The framework provides a way to explain how humans interpret events -positive, negative- as well as their position, support and engagement with respect to their own interpretation and report – confidence, support, agreement, disagreement, as shown in tab. I.

From a linguistic standpoint, the appraisal theory describes how authors use linguistic expressions to communicate their emotional states and engagement.

TABLE I
EXAMPLE OF APPRAISAL CATEGORIES SPLIT INTO POLARITY

Polarity	Appreciation	Engagement	Affect
Positive	Amazing	Obviously	Joy, happy
Negative	Awful	Apparently	Miserable

The appraisal theory structures appraisal expressions under three main basic categories describing attitudes, engagement and graduation, as shown in fig. 2.




Fig. 2. Appraisal categories

The attitude system is related to linguistic expressions conveying the attitude of authors at the time they write the text. This system covers three main subcategories: affect, appreciation and judgement, discussed hereafter.

Affect is related to linguistic expressions of author's feelings such as happiness, joy, sadness, grief, etc.

Judgment highlights linguistic expressions conveying characterization of persons and behaviors by the author. Generally it conveys opinions and personal tastes about objects, such as nice, ugly, beautiful, shy but also about interactions and behaviors in the social context: heroic, brave, open-minded.

Appreciation is related to assessment and evaluations of entities, objects, events and scenes.

The *engagement system* gathers linguistic expression specifying the author's position with respect to his own statements. When reporting, writers often embed clues as to how strongly they support the content being conveyed and may indicate confidence, doubt, skepticism, conviction, etc., about the information reported. The engagement system is closely related to the notions of trust, confidence, probability or possibility.

Categories under this system encompass aspects related to denial, concession, confirmation, endorsement, acknowledgement and distance.

Based on this frame, the appraisal ontology [43] was built in order to have a general description of appraisal categories and to provide a formal model of concepts and their relationship.

The ontology was created from scratch, starting with categories introduced by the appraisal theory, and adding additional concepts to characterise appraisal expressions.

Namely, *graduation* concept was enriched by adding several subclasses to describe Focus as sharp or soft and Force as high or low. Under graduation systems as well, Polarity is a concept used to capture positive and negative emotional states conveyed by appraisal expressions.

The appraisal ontology is represented in OWL [44] and is composed of 46 concepts structured on a 6-levels hierarchy; the model also has 4 ObjectProperties and 2 DataType propertie, 50 nodes, 75 edges and 268 instances of concepts.

By using appraisal categories, the annotation goes beyond limitations of traditional concepts of sentiment and opinions and highlights the set of appraisal expression, which is a linguistic unit by which a personal appreciation is conveyed, whether it is an opinion, sentiment or supportive statements.

IV. EXPERIMENTS

Experiments were performed to evaluate the performance of the overall trend detection approach.

A. Experimental settings

Datasets: Data acquisition was done by crawling several sets of tweets with specific keywords and additional constraints to select only posts written in English. This phase was carried out as a straightforward step and relies only on data content of both keywords and posts and the ability of API used to search the social network. Three distinct data sets were collected based on query terms, without manual processing or tagging:

- Data set 1: *Brexit*
- Data set 2: *Western values, White supremacy*
- Data set 3: *Security, Hate occident*

Preprocessing: This step involves a series of techniques aiming to optimize the performance of experiments. Prior to conducting experiments, tweets are broken down into sentences and then the preprocessing includes: tokenization, lower casing all tokens, removing images, punctuation and stop words. Furthermore usernames, retweets, duplicates and web references are removed as well.

Tab. II shows the number of tweets before and after the preprocessing step and also highlights a significant drop in tweets after the preprocessing.

TABLE II
DATA SETS

Number of	Data set 1	Data set 2	Data set 3
Original tweets	103 764	51 000	320 000
Preprocessed tweets	12 057	64 203	124 384

The rationale behind this cleaning step is that tweets often get large number of retweets and keeping them in the dataset will produce large clusters containing redundant tweets, and it will reinforce false patterns.

Numerical representation : The initial word representations of tweets are initialized with a 300 dimension GloVe

vectors which are pretrained [45]. Data sets are split into three parts, with 60/20/20 ratios for training, validation and test sets. Learning algorithms were implemented and run in Python.

Semantic annotation for appraisal detection: attaches additional information to tweets based on their content. Semantic annotation is performed automatically, by using lexical similarities measures that associate a real number to a pair of words and offers a measure of the degree to which two words are similar. For this work, lexical similarities are used to assign words to concepts of the appraisal ontology.

Experimental protocol: Each tweet is treated as an individual document for classification and appraisal annotation purpose. After classification, tweets belonging to the same cluster are combined into one document to both extract dominant topics with LDA and aggregate individual appraisal categories into a general trend at cluster level. The reason for pooling tweets into one document is to overcome the sparsity of twitter data for specific topics or appraisal categories.

B. Results and evaluation

Results of clustering and topics detection : Parameters for unsupervised clustering were tuned by using the Elbow heuristic method. This method takes into account the variance, explained as a function of the number of clusters to be build and identifies graphically an elbow threshold highlighting to the number of clusters corresponding to the optimal value of variance.

Mathematically,

$$W_k = \sum_{r=1}^k D_r / n_r \quad (1)$$

Where k is the number of clusters, n_r is the number of points in cluster r and D_r is the sum of distances between all points in a cluster:

$$D_r = \sum_{i=1}^{n_r-1} \sum_{j=1}^{n_r} ||d_i - d_j||^2 \quad (2)$$

Tab. III shows the number of clusters created according to the Elbow criterion and the percentage of clustered data for each data set.

TABLE III
NUMBER OF CLUSTERS

	Data set 1	Data set 2	Data set 3
Number of clusters	2	5	10
Clustered data (%)	97	81	76

In addition, topics coherence scores were used to evaluate the coherence of clusters and the results were visualized by using the pyLDavis library ¹. Intuitively, the coherence measure captures the precision of the clustering, and values under 100 are due to the fact that the categorization of tweets is not crisp, and some tweets belong to several clusters.

¹<https://pypi.org/project/pyLDavis/>

For the first data set, the clustering algorithm identified 2 clusters illustrated in fig. 3, and topic detection highlighted several words associated to each cluster.




Fig. 3. Clusters of data set 1 (3D)

Among the topics, a main topic clearly highlights Theresa May as dominant named entity.

The number of clusters for the second data set is 5.

TABLE IV
EXAMPLE OF TOPICS AND ITS MOST FREQUENT WORDS

topic 1	supremacy	commoner	sacrifice	opinion
topic 2	protect	surprise	person	promote
topic 3	patriarchy	racism	child	never

Tab. IV shows a topic as build by LDA. Topics detection clearly identifies topics related to white supremacy, although the overall data set was collected with 2 keywords.

The processing of the last set provides 10 clusters, see fig. 4. This data set is the most heterogeneous, due to the keyword security that gathers a large variety of tweets.




Fig. 4. Clusters for data set 3

The distribution of topics detected within is also sparse, as illustrated in fig. 5.

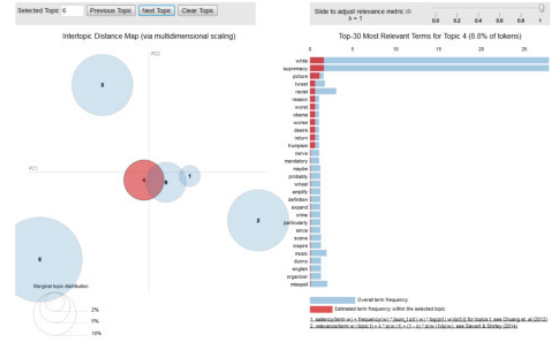


Fig. 5. Visualization of association topics-clusters

Appraisal annotation : After the clustering phase, the appraisal ontology was used to detect instances of appraisal categories at tweet level. Thus, tweets were annotated through the ontology of appraisals. Tab. V shows the analysis of data sets in terms of numerical distribution of subjective and objective tweets.

TABLE V
PERCENTAGE OF ANNOTATIONS

	Data set 1	Data set 2	Data set 3
Annotated data (%)	96	83	79

Numerical values show a high percentage of subjective tweets for all datasets analyzed, regardless of the keyword used to collect. The trend is still the same at cluster level, with clusters having similar values for the ration of subjective/objective content.

Fig. 6 illustrates the distribution of main appraisal categories for the three data sets.




Fig. 6. Distribution of main appraisal categories in data sets

The main appraisal categories detected are: *Attitude* (as a general concept), *affect* (specific concept under attitude), and *support*, *deny*, *disapproval* and *submission*, all specific concepts under *Engagement* concept.

Thus, data set1 collected with Brexit keyword, appears related to *affect*, *deny*, *disapproval* and *submission* concepts while being less correlated with concepts under *Attitude* and

Support. The same associations appear when analyzing individually each cluster.

Data set 2 is strongly correlated with *Deny* concept, and has an almost equivalent distribution of *affect*, *attitude*, *support*, *submission* and *disapproval* categories.

Data set 3 has a strong affect component, and is strongly correlated with *Attitude* and *Affect*. Under *Engagement* system, *Disapproval* is the most relevant concept associated to this data set.




Fig. 7. Distribution of polarities in data sets

Fig. 7 shows the distribution of positive and negative orientation, with high and low force.

Regarding the orientation of tweets, there is a good representation of negative-oriented tweets that account for more than 50% for each collection. Although the negative orientation is prevalent, the majority of tweets are of moderate impact, and tweets having low or high impact are underrepresented for all data sets.

Evaluation of trend detection Quantitative analysis by manually inspecting the quality of output was also undertaken to validate the results provided by the overall approach. Ground truth was created for the data set 1 by analyzing the type and polarity of annotation in the light of Google news headlines, see tab. VI.

TABLE VI
MANUAL VALIDATION AND GROUND TRUTH IN %

True positives (Tp)	7211
Trues negatives (Tn)	972
False positives (Fp)	913
False negatives (Fn)	2300
Undetermined (U)	661

In the light of those values, the quality of trend detection is estimated with *Precision*, *Recall* and *Accuracy*. Values of metrics (%) are: Precision = 88; Recall = 75, Accuracy = 72.

Statistical values discussed above give a general overview of the association between data sets, appraisal and their orientations as build thanks to the clustering algorithms and semantic annotation. In addition, a manual analysis of results sheds light on a number of findings that can be useful in better understanding such diverse associations of subjective states and real-life topics.

Thus, through the appraisal pooling strategy, positive polarities are associated with Theresa May, whilst negative categories often included references to political system, recession and war. Similarly, Brexit was frequently mentioned in tweets associated later to topics that belonged to positive polarities, including terms such opinion, person and symbol.

From a different perspective, topics such as *supremacy* were associated with negative emotions such as deny and disapproval.

And finally, within the last data set, terms such as protect, surprise and patriarchy were closely associated to categories such as submission, deny and affect.

V. DISCUSSION

The main purpose of this work was to develop a hybrid approach to detect trends with regards to topics on social media. The work was done under the hypothesis that clustering data can pool into the same cluster similar tweets and the appraisal categories are able to highlight words that are more frequently used when people hold a specific perspective. Results are impacted by several factors discussed hereafter.

Data sets and clustering techniques

Data sets collected for this work are relatively small, especially as a result of the preprocessing step. From a practical standpoint, adding the unsupervised clustering step as such does not offer new insights into trend analysis, since the association of appraisal categories are very similar from one cluster to another. Topic detection, instead, offers a good basis to analyze the results and shed light on words that are relevant to describe the clusters. However, clustering algorithms are based on numerical representations of textual data, and thus ignores the order of appearance of words and the correlations between words.

Coverage of the appraisal ontology

The appraisal ontology was build to capture features of subjectivity according to principles of the appraisal theory. The terminological dimension is emphasized by adding instances of concepts, and also by highlighting how a broader or narrower meaning can be created by linguistic adjustments. Moreover, elements of the model are drawn from existing thesauri and semantic lexicons.

Impact of sarcasm and irony

Negative of positive polarity is associated to concept of the appraisal ontology, whilst the polarity of words can change according to the context or due to sarcasm and irony. The accepted view is that sarcasm and irony potentially flip the polarity of tweets. As polarity is modeled as a concept of the ontology having three exclusive values (positive, negative and neutral), overcoming limitations due to ambiguous polarity requires the implementation external procedures, able to detect the occurrence context of words and infer their polarity accordingly.

VI. CONCLUSION AND FUTURE WORK

This paper presents a hybrid approach to detect trends in social data. Trends are understood as associations of topics and appraisal categories, along with their positive and negative polarity. The solution first uses unsupervised clustering to pool data into more homogeneous clusters; topics are that identified to provide richer description of clusters by highlighting specific words; in parallel, data is also annotated with concepts of the appraisal ontology and the content is described in terms of affects, attitude, judgment, deny, engagement, disapproval or rejection. Experiments have been conducted with three data sets and results are partially analyzed against manually crafted ground truth.

Directions for future work are threefold. First, more features should be considered for unsupervised classification such as the hashtags of tweets, in order to improve the coherence and quality of clusters and boost performance.

Second, the analysis of polarity should be extended by considering not only the annotation given by the appraisal ontology, but also the context of appearance of words, as to acknowledge and take into account the impact of linguistic phenomena.

And finally, since the manual analysis of results for validation is time consuming and prone to errors and bias, a new research direction would be to implement the evaluation procedures by comparing the outcome with results provided by using different resources, such as SenticNet [46].

Acknowledgement

Results reported in this document have been produced in the context of the PRACTICES project. Research has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 740072.

REFERENCES

- [1] P. Barberá, A. Casas, J. Nagler, P. J. Egan, R. Bonneau, J. T. Jost, and J. A. Tucker, "Who leads? who follows? measuring issue attention and agenda setting by legislators and the mass public using social media data," *American Political Science Review*, vol. 113, no. 4, pp. 883–901, 2019.
- [2] A. Sheth, P. Anantharam, and C. Henson, "Physical-cyber-social computing: An early 21st century approach," *IEEE Intelligent Systems*, no. 1, pp. 78–82, 2013.
- [3] P. Wang, R. Angarita, and I. Renne, "Is this the era of misinformation yet: combining social bots and fake news to deceive the masses," in *Companion Proceedings of the The Web Conference 2018*. International World Wide Web Conferences Steering Committee, 2018, pp. 1557–1561.
- [4] M. J. Mazarr, A. Casey, A. Demus, S. W. Harold, L. J. Matthews, N. Beauchamp-Mustafaga, and J. Sladden, "Hostile social manipulation present realities and emerging trends," RAND Corporation Santa Monica United States, Tech. Rep., 2019.
- [5] J. Jouhki, E. Lauk, M. Penttinen, N. Sormanen, and T. Uskali, "Facebook's emotional contagion experiment as a challenge to research ethics," *Media and Communication*, vol. 4, 2016.
- [6] E. Ferrara, "Contagion dynamics of extremist propaganda in social networks," *Information Sciences*, vol. 418, pp. 1–12, 2017.
- [7] S. M. Mohammad, "Sentiment analysis: Detecting valence, emotions, and other affectual states from text," in *Emotion measurement*. Elsevier, 2016, pp. 201–237.
- [8] S. Mohammad, F. Bravo-Marquez, M. Salameh, and S. Kiritchenko, "Semeval-2018 task 1: Affect in tweets," in *Proceedings of the 12th international workshop on semantic evaluation*, 2018, pp. 1–17.
- [9] C. Vicient and A. Moreno, "Unsupervised topic discovery in microblogging networks," *Expert Systems with Applications*, vol. 42, no. 17–18, pp. 6472–6485, 2015.
- [10] B. Wang, M. Liakata, A. Zubiaga, and R. Procter, "A hierarchical topic modelling approach for tweet clustering," in *International Conference on Social Informatics*. Springer, 2017, pp. 378–390.
- [11] J. Zeng, J. Li, Y. He, C. Gao, M. R. Lyu, and I. King, "What you say and how you say it: Joint modeling of topics and discourse in microblog conversations," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 267–281, 2019.
- [12] P. Sriji, M. Hepple, K. Bontcheva, and D. Preotiu-Pietro, "Sub-story detection in twitter with hierarchical dirichlet processes," *Information Processing & Management*, vol. 53, no. 4, pp. 989–1003, 2017.
- [13] A. C. E. Lima, L. N. de Castro, and J. M. Corchado, "A polarity analysis framework for twitter messages," *Applied Mathematics and Computation*, vol. 270, pp. 756–767, 2015.
- [14] A. Giachanou and F. Crestani, "Like it or not: A survey of twitter sentiment analysis methods," *ACM Computing Surveys (CSUR)*, vol. 49, no. 2, p. 28, 2016.
- [15] K. Ravi and V. Ravi, "A survey on opinion mining and sentiment analysis: tasks, approaches and applications," *Knowledge-Based Systems*, vol. 89, pp. 14–46, 2015.
- [16] Q. Zhou and C. Zhang, "Emotion evolutions of sub-topics about popular events on microblogs," *The Electronic Library*, vol. 35, no. 4, pp. 770–782, 2017.
- [17] Y. Hu, X. Xu, and L. Li, "Analyzing topic-sentiment and topic evolution over time from social media," in *International conference on knowledge science, engineering and management*. Springer, 2016, pp. 97–109.
- [18] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, "DbSCAN revisited, revisited: why and how you should (still) use dbSCAN," *ACM Transactions on Database Systems (TODS)*, vol. 42, no. 3, p. 19, 2017.
- [19] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [20] H. Alvari, S. Sarkar, and P. Shakaran, "Detection of violent extremists in social media," *arXiv preprint arXiv:1902.01577*, 2019.
- [21] S. Kannangara, "Mining twitter for fine-grained political opinion polarity classification, ideology detection and sarcasm detection," in *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 2018, pp. 751–752.
- [22] T. P. Sahu and S. Ahuja, "Sentiment analysis of movie reviews: A study on feature selection and classification algorithms," in *2016 International Conference on Microelectronics, Computing and Communications (MicroCom)*. IEEE, 2016, pp. 1–6.
- [23] B. Liu, "Sentiment analysis and opinion mining," *Synthesis lectures on human language technologies*, vol. 5, no. 1, pp. 1–167, 2012.
- [24] J. Deriu, A. Lucchi, V. De Luca, A. Severin, S. Müller, M. Cieliebak, T. Hofmann, and M. Jaggi, "Leveraging large amounts of weakly supervised data for multi-language sentiment classification," in *Proceedings of the 26th international conference on world wide web*, 2017, pp. 1045–1052.
- [25] D. Tang, B. Qin, and T. Liu, "Learning semantic representations of users and products for document level sentiment classification," in *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2015, pp. 1014–1023.
- [26] M. d. P. Salas-Zarate, R. Valencia-Garcia, A. Ruiz-Martinez, and R. Colomo-Palacios, "Feature-based opinion mining in financial news: an ontology-driven approach," *Journal of Information Science*, vol. 43, no. 4, pp. 458–479, 2017.
- [27] K. Chakraborty, S. Bhattacharyya, R. Bag, and A. E. Hassanien, "Comparative sentiment analysis on a set of movie reviews using deep learning approach," in *International Conference on Advanced Machine Learning Technologies and Applications*. Springer, 2018, pp. 311–318.
- [28] F. Hemmatian and M. K. Sohrabi, "A survey on classification techniques for opinion mining and sentiment analysis," *Artificial Intelligence Review*, pp. 1–51, 2017.
- [29] L. Chen, J. Martineau, D. Cheng, and A. Sheth, "Clustering for simultaneous extraction of aspects and features from reviews," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016, pp. 789–799.
- [30] M. Thelwall and B. Cugelman, "Monitoring twitter strategies to discover resonating topics: The case of the undp," 2017.

- [31] Z. Li, Y. Zhang, Y. Wei, Y. Wu, and Q. Yang, "End-to-end adversarial memory network for cross-domain sentiment classification." in *IJCAI*, 2017, pp. 2237–2243.
- [32] F. H. Khan, U. Qamar, and S. Bashir, "Sentimi: Introducing point-wise mutual information with sentiwordnet to improve sentiment polarity detection," *Applied Soft Computing*, vol. 39, pp. 140–153, 2016.
- [33] A. Agarwal, V. Sharma, G. Sikka, and R. Dhir, "Opinion mining of news headlines using sentiwordnet," in *2016 Symposium on Colossal Data Analysis and Networking (CDAN)*. IEEE, 2016, pp. 1–5.
- [34] A. Esuli, "The user feedback on sentiwordnet," *arXiv preprint arXiv:1306.1343*, 2013.
- [35] A. Schoene and G. de Mel, "Pooling tweets by fine-grained emotions to uncover topic trends in social media."
- [36] P. Vijayaraghavan, S. Vosoughi, and D. Roy, "Automatic detection and categorization of election-related tweets," in *Tenth International AAAI Conference on Web and Social Media*, 2016.
- [37] D. Zimbra, A. Abbasi, D. Zeng, and H. Chen, "The state-of-the-art in twitter sentiment analysis: A review and benchmark evaluation," *ACM Transactions on Management Information Systems (TMIS)*, vol. 9, no. 2, pp. 1–29, 2018.
- [38] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise." in *Kdd*, vol. 96, no. 34, 1996, pp. 226–231.
- [39] G. J. King, R. R. Richards, J. D. Zuckerman, R. Blasier, C. Dillman, R. J. Friedman, G. M. Gartsman, J. P. Iannotti, J. P. Murnahan, V. C. Mow *et al.*, "A standardized method for assessment of elbow function," *Journal of shoulder and elbow surgery*, vol. 8, no. 4, pp. 351–354, 1999.
- [40] R. Harakawa, S. Takimura, T. Ogawa, M. Haseyama, and M. Iwahashi, "Consensus clustering of tweet networks via semantic and sentiment similarity estimation," *IEEE Access*, vol. 7, pp. 116 207–116 217, 2019.
- [41] J. Boyd-Graber and P. Resnik, "Holistic sentiment analysis across languages: Multilingual supervised latent dirichlet allocation," in *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2010, pp. 45–55.
- [42] P. R. White, "Appraisal theory," *The international encyclopedia of language and social interaction*, pp. 1–7, 2015.
- [43] V. Dragos, D. Battistelli, and E. Kelodjoue, "Beyond sentiments and opinions: exploring social media with appraisal categories," in *2018 21st International Conference on Information Fusion (FUSION)*. IEEE, 2018, pp. 1851–1858.
- [44] D. L. McGuinness, F. Van Harmelen *et al.*, "Owl web ontology language overview," *W3C recommendation*, vol. 10, no. 10, p. 2004, 2004.
- [45] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [46] E. Cambria, S. Poria, R. Bajpai, and B. Schuller, "Senticnet 4: A semantic resource for sentiment analysis based on conceptual primitives," in *Proceedings of COLING 2016, the 26th international conference on computational linguistics: Technical papers*, 2016, pp. 2666–2677.