

Regrexit or not Regrexit: Aspect-based Sentiment Analysis in Polarized Contexts

Vorakit Vorakitphan¹, Marco Guerini², Elena Cabrio³, Serena Villata³

Université Côte d’Azur, CNRS, Inria, I3S, France^{1,3}

Fondazione Bruno Kessler, Trento, Italy²

vorakit.vorakitphan@inria.fr¹, guerini@fbk.eu²,
{firstname.lastname}@univ-cotedazur.fr³

Abstract

Emotion analysis in polarized contexts represents a challenge for Natural Language Processing modeling. As a step in the aforementioned direction, we present a methodology to extend the task of Aspect-based Sentiment Analysis (ABSA) toward the affect and emotion representation in polarized settings. In particular, we adopt the three-dimensional model of affect based on Valence, Arousal, and Dominance (VAD). We then present a Brexit scenario that proves how affect varies toward the same aspect when politically polarized stances are presented. Our approach captures aspect-based polarization from newspapers regarding the Brexit scenario of 1.2m entities at sentence-level. We demonstrate how basic constituents of emotions can be mapped to the VAD model, along with their interactions respecting the polarized context in ABSA settings using biased key-concepts (e.g., “stop Brexit” vs. “support Brexit”). Quite intriguingly, the framework achieves to produce coherent aspect evidences of Brexit’s stance from key-concepts, showing that VAD influence the support and opposition aspects.

1 Introduction

Aspect-based Sentiment Analysis (ABSA) aims at capturing sentiment (i.e., positive, negative or neutral) expressed toward each aspect (i.e., attribute) of a target entity. The main interest is to capture sentiment nuances about different entities. However, in a context of opinion polarization, different groups of people can form strong convictions of competing opinions on such target entities, resulting in different (often opposite) evaluations of the same aspect. Compare, for example, the differences in the pro- and anti-Brexit discourses concerning the withdrawal of the United Kingdom from the European Union (EU), aligning with contrasting attitudes toward the EU, the immigration and the country’s culture. In fact, while in standard scenarios of sentiment analysis about specific entities and their aspects it is assumed that sentiment is consistent (e.g., *a big screen is a desirable characteristics for a TV*), this is not the case for polarized contexts. Hence, for example, a “clean Brexit” might be desirable to some, but not to others.

Some previous work has showed that sentiment alone (i.e., positive, negative, neutral) is not able to grasp complex emotion phenomena in information diffusion scenarios (Guerini and Staiano, 2015; Berger and Milkman, 2012). Following the suggestion of (Guerini and Staiano, 2015), to analyse in a fine-grained manner our scenario where the same topic can appear in different news outlets that have opposed opinions, we decide to make resort to the model presented in (Russell, 1980; Bradley and Lang, 1994). This work developed a circumplex model, with a three dimensional space of affect along the Valence, Arousal, and Dominance coordinates, where all emotions can be mapped in the same dimensional space. In our work, we therefore hypothesize that VAD plays a significant role in the interaction with linguistic elements in ABSA settings where polarized opinions are at play.

More specifically, we answer the following research questions: *how basic constituents of emotions, such as VAD, interact with aspect-based polarization? How polarized settings differ from standard SA scenarios? How aspects of Brexit and related key-concepts are perceived by opposite parties?*

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

To answer these research questions, we propose a comprehensive framework for studying the interaction of ABSA with opinion polarization in newspapers and social media. We first trained an RNN that detects the emotions and their intensities at sentence-level, and then we map emotion intensities into the VAD circumplex model. Later, we build a framework to assess whether and how VAD are connected to polarized contexts, by computing the VAD scores of a set of key-concepts that can be found on newspapers with opposite views. These key-concepts (e.g., “stop immigration”) are built from a set of aspects (“immigration” in our example) combined with relevant verbs or adjectives that represent a clear polarized opinion toward the aspect (e.g., “stop”).

To experiment with the proposed approach, we focus on the Brexit scenario, whose unique circumstances have boosted opinion polarization because of the extremely rapid and volatile changes in the political panorama as the results of the EU referendum were announced (political leaders were quick to say that the outcome had “revealed a divided Britain”)¹. While an analysis of this situation from the viewpoint of political sciences is not the focus of our work, the Brexit scenario as discussed on newspapers and on social media provides us with the required elements to carry out our study, because of the opinion divisions formed around one or more political positions or issues. Moreover, to be meaningfully polarizing, issues need to be important (or “salient”) to a large section of the public, not just a minority of people with strongly held views, and in the Brexit scenario this is actually the case. In our experimental setting, we select two British newspapers known to be polarized, i.e., either for or against Brexit. Results show that VAD are not absolute, but relative to the newspaper’s viewpoint on the key-concept. Our approach highlights that using the proposed key-concepts gives us fine-grained details about VAD elements that strongly interact with the polarized context. We show that standard SA approaches can be deceptive in such polarized setting (considering only the word “Brexit” on both newspapers, the valence is almost identical), while our ABSA approach shows a clear-cut polarization.

2 Related Work

ABSA approaches aim not only at capturing linguistic relations, but also at ascribing polarity stances of textual entities. We concentrate our attention to this second aspect being it closer to the purpose of our work. Demszky et al. (2019) apply ABSA methods to pioneer tweets on the US mass shootings topic, where the topic was politically discussed from different viewpoints according to the locations of events with the contrasting use of the terms “terrorist” and “crazy”, that contribute to polarization. Still, while in (Demszky et al., 2019) such terms were taken as evidence of polarization, in our case we want to understand how such concepts are evaluated – in polarized scenarios – by both parties (e.g., how people and journalists from each stance react when the same shooter is defined as “terrorist” or “crazy”). Other approaches in this direction employ NLP to show that certain topics are more polarizing than others (Balasubramanyan et al., 2012). Moreover, (Zhang et al., 2019) and (Chen et al., 2020) show that most recent ABSA approaches rely on state-of-the-art supervised approaches combining complex layers of neural network models (e.g., transformers) to classify labels representing aspects from text elements with standard sentiment (i.e., positive, negative). This direction is computationally expensive and requires the annotation of the aspects. In contrast, our ABSA method consists in an unsupervised approach where aspects are conveyed with respect to fine-grained extraction of VAD affects in polarized situations.

3 Emotions Analysis

As the first step of the proposed framework, we focus on the task of detecting emotions and their intensities at sentence-level from the news context. We compare two standard emotion recognition approaches, i.e., a lexicon-based model called DepecheMood++ (Araque et al., 2019), and an RNN based approach that casts the task of multi-label ABSA as a regression problem. Both DepecheMood++ and the RNN model were built and trained on the same set of news extracted from the *Rappler.com* website. *Rappler.com* is an online news publisher, where each article is associated to a mood meter. Such mood meter allows each reader to vote the emotions evoked by the article after reading. We implemented a

¹<https://www.kcl.ac.uk/policy-institute/assets/divided-britain.pdf>

Emotions	DepecheMood++	LSTM	BiLSTM	CNN
ANGER	0.47	0.62	0.67	0.64
FEAR	0.60	0.78	0.76	0.73
JOY	0.38	0.64	0.63	0.51
SADNESS	0.46	0.75	0.75	0.7
SURPRISE	0.21	0.43	0.52	0.48
AVG All Emotions	0.43	0.64	0.66	0.61

Table 1: Pearson’s correlation (r) on Cross-dataset Testing of Emotion Recognition Models.

web crawler and harvested a total of 67,828 articles from January 2017 until April 2019 along with the 8 possible emotions voted by readers (i.e., *happy, inspired, amused, afraid, angry, annoyed, don’t care, sad*). As we aim at detecting emotions at sentence-level, we pick all headlines of such articles and the associated emotions as our training set. For the RNN based approach, we use a deep contextualized word representation as our features, namely ELMo (Peters et al., 2018), with a deep bidirectional language model (BiLSTM). We use BiLSTM of 128 dimensional layers with a Rectified Linear Unit (ReLU) activation function (Agarap, 2018). Then, we apply Mean Squared Error (MSE) as a loss function with a standard adam optimizer. At the last layer, we use sigmoid function to gate our multi-label output neurons as the intensity of each emotion. To validate our model, we perform cross-dataset testing on the “Affective Text” dataset from SemEval2007-T14 (Strapparava and Mihalcea, 2007) using Pearson’s correlation (r) to evaluate the performances of our system on each individual emotion. In Table 1², we report the results for each emotion, and the average of all emotions produced by the different models. Given that the ELMo embeddings with BiLSTM model outperformed the other models, we selected the output of this model as input for the next step (i.e., to map the predicted emotions to the VAD model).

4 VAD Analysis

The VAD model of affects (Russell, 1980; Bradley and Lang, 1994) has been extensively used, and relies on three dimensions: *Valence*, that can be positive/negative (e.g., FEAR has a negative valence, JOY has a positive valence) and corresponds to the standard dimension of sentiment analysis; *Arousal*, that can be low or high and ranges from “calm” to “excitement” (e.g., SADNESS has low arousal while ANGER has high arousal, even if they have the same sentiment); finally, *Dominance*, that can be low or high as well, and ranges from “controlled” to “in control” (e.g. FEAR refers to low dominance, INSPIRED to high dominance).

To map the emotions expressed in a sentence to the VAD model, we followed the approach of Guerini and Staiano (2015). We use the VAD scores of emotion labels provided in (Warriner et al., 2013) which serves as our gold standard to obtain VAD scores from emotion intensities. In particular, here below we report the procedure for Valence computation (Arousal and Dominance computation are akin): given a sentence S , to obtain its valence S_v we multiply the intensity of each emotion $I(e_i)$ – computed using the approach described in Section 3 – with the corresponding Valence score $V(e_i)$ present in (Warriner et al., 2013). Then, we sum the n emotional dimensions to obtain the final Valence score, as expressed in Equation 1:

$$S_v = \sum_{i=1}^n V(e_i) \times I(e_i) \quad (1)$$

After the conversion, VAD values are in absolute form (arbitrary range). Since newspapers could have different “affective styles” – that are prior to their stance on a specific topic – we apply standardization individually to each VAD dimension for each newspaper to make results better comparable across newspapers.

²The “DISGUST” label is excluded as there is no correspondence to Rappler emotions.

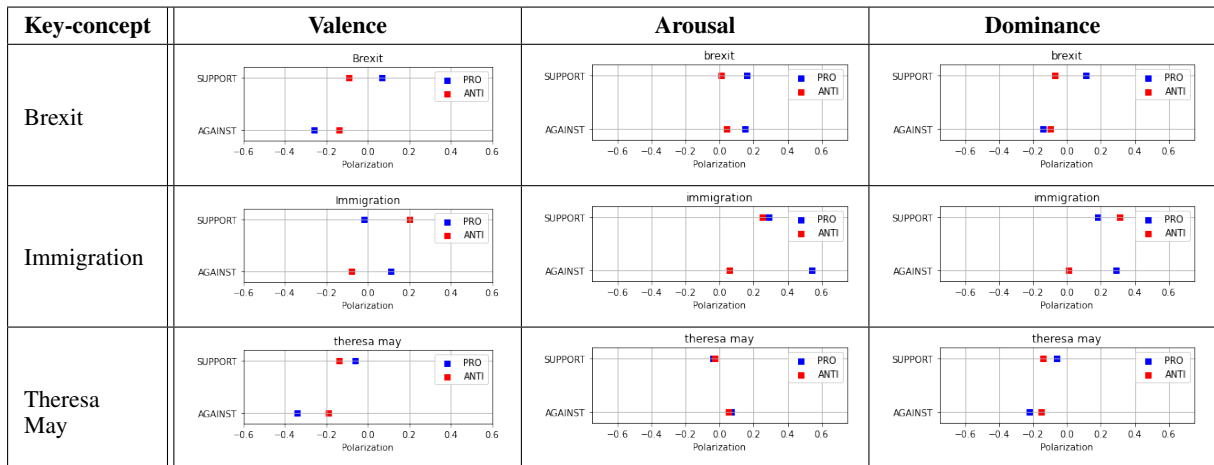


Table 2: Polarization on VAD Dimensions using Key-Concepts Approach.

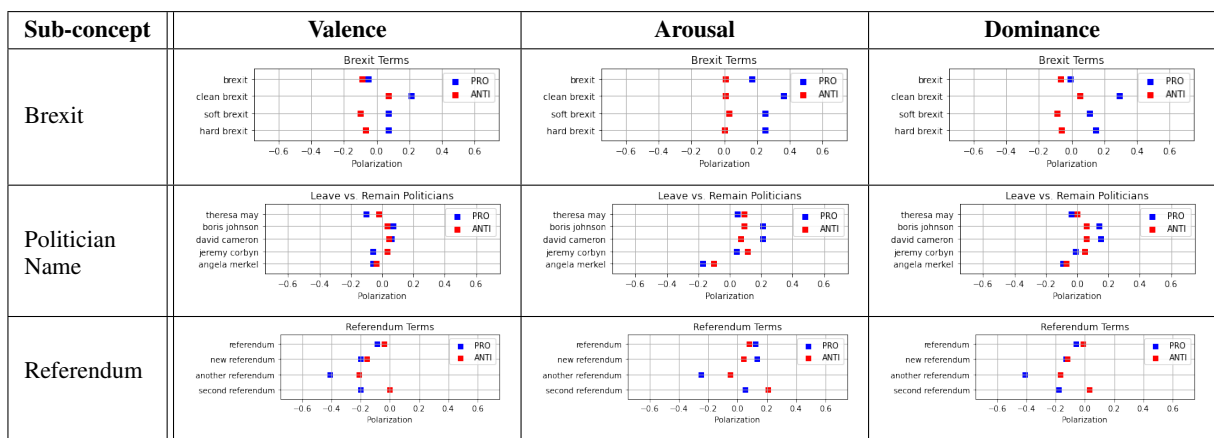


Table 3: Polarization on VAD Dimensions on Sub-concepts via Averaging Topic.

5 A Polarized Context Scenario: the Brexit

Dataset collection. To clearly define the use case, we conducted a survey by asking to 10 British participants currently living in England to agree on the stance of UK newspapers with respect to the Brexit, i.e., they were asked which newspapers are known to be pro- or anti- Brexit. As a result, we selected two newspapers that obtained the mutual agreement across all participants, i.e., The Sun (*thesun.co.uk*) as the pro-Brexit newspaper, and The Guardian (*theguardian.com*) as the anti-Brexit one. We built a web crawler to harvest the articles containing the word “brexit” from these two news websites from year 2017 to 2019. We obtained 28,212 articles in total. We segmented them at sentence-level resulting in 1.2 million sentences.

Key-concepts. For our analysis, we focused on a set of key-concepts relevant to the Brexit discussion (e.g., “brexit”, “immigration”, “EU”). We created two lists of words that could potentially trigger polarization occurring with the identified key-concepts in text (e.g., “stop”, “promote”). These words have then been combined with each key-concept to obtain the “support” and “against” aspect clusters, e.g., “stop Brexit”, “block Brexit” for the against aspect, and “support Brexit”, “make Brexit” for the support aspect. We then computed the VAD scores for the sentences in our dataset containing an occurrence of the mentioned aspects. We used weighted average on all sentences collected for each aspect cluster to compute the final VAD scores for the support and against aspect clusters. For consistency reasons, we select key-concepts only when their frequency (in both support and against aspect clusters) > 1,000.

Results and discussion. Table 2 shows some clear-cut results for key-concepts “Brexit”, “Immigration” and “Theresa May”. Considering “Brexit” key-concept, we see that the valence of support aspect in

pro-Brexit journal is positive, while it is negative in anti-Brexit (as expected). Turning to arousal, we see that the concept “Brexit” provokes a higher emotional activation in pro-Brexit journal, regardless of the aspect being support or against. “Immigration”, instead, is the key-concept with the highest arousal in all the considered cases, both for pro- and for anti- Brexit in the support case, but with opposite valence (positive for anti and neutral/negative for pro). In the against aspect instead (e.g., “stop immigration”), pro-Brexit have an extreme arousal connected to positive sentiment (valence), while anti-Brexit have low arousal with slightly negative sentiment (i.e., a resignation attitude). Interestingly, dominance (the feeling of being in control of the situation) is consistent with the polarized sentiment in both aspects and journals, with the “resignation attitude” confirmed also in this case for the against aspect of anti-Brexit. Lastly, we consider a politician, “Theresa May”, giving that at the time we crawled the articles she was the UK Prime Minister. Results show a convergence toward a neutral score on arousal for both pro- and anti-Brexit (both for support and against key-concepts). For what concerns valence and dominance, they are both in the negative area of the polarization: the anti-Brexit journal has a more negative valence in the cluster of the support aspect (i.e., they are against Theresa May), while the opposite holds for the against aspect cluster, where the pro-Brexit journal has, as expected, more negative valence. These results seem to imply that the pro-Brexit journal has a “lukewarm attitude” toward Theresa May (neutral arousal, almost neutral sentiment in support but negative sentiment when she is criticized).

Finally, in Table 3, we analysed some concepts for which we did not have enough occurrences to create the support and against aspect clusters, but are nonetheless important with respect to the journal stance. Referring to sub-concepts of the Brexit scenario, i.e., “clean Brexit”, “soft Brexit” and “hard Brexit”, we see a clear consistency on all VAD dimensions with the previous analysis. It is also interesting to note that without using the proposed aspects, we obtain a neutral position on the Brexit key-concept alone, and inexact positions on Theresa May as well (compare the two key-concepts in Tables 2 and 3). This proves that in polarized settings, applying just standard SA to perform ABSA can be misleading for concepts with high occurrences, which is mostly the case in traditional sentiment extraction. Other sub-concepts of Brexit, such as “clean Brexit”, “soft Brexit” and “hard Brexit”, have highlighted the consistency of VAD trends such that the pro-Brexit journal reveals higher scores on all VAD elements for these sub-concepts, than the anti-Brexit one. The same holds for sub-concepts concerning “politician name” (i.e., pro-Brexit politicians have a positive valence in pro-Brexit journals and anti-Brexit politicians have better valence in anti-Brexit journals). Finally, also for “referendum” anti-Brexit journal shows a consistently higher valence (i.e., preference toward an additional referendum).

6 Conclusions and Future Work

In this work, we presented a methodology to extend the task of Aspect-based Sentiment Analysis so that to include affect and emotions representation in polarized settings. In particular, we adopted the three-dimensional model of affect based on Valence, Arousal, and Dominance. We tested the proposed framework on the Brexit scenario, showing how affect dimensions (VAD) vary on the same aspect when politically polarized stances are presented. Our approach was able to capture stereotypical aspect-based polarization from newspapers regarding the Brexit scenario using biased key-concepts (e.g., “stop Brexit” vs. “support Brexit”).

As for future work, we plan to investigate the correlation of our results with virality indexes on social media (e.g., number of likes and retweets, number of replied messages). We also plan to have a comparative analysis with the textual material produced by the large audience that is supposed to share the same stance/polarization as the newspaper considering the comments to the articles.

Acknowledgement

This work is partially supported by the ANSWER project PIA FSN2 n. P159564-2661789/DOS0060094 between Inria and Qwant. This work has been supported by the French government, through the 3IA Côte d’Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR- 19-P3IA-0002.

References

- Abien Fred Agarap. 2018. Deep Learning using Rectified Linear Units (ReLU). arXiv e-prints, page arXiv:1803.08375, March.
- O. Araque, L. Gatti, J. Staiano, and M. Guerini. 2019. Depechemood++: a bilingual emotion lexicon built through simple yet powerful techniques. IEEE Transactions on Affective Computing, pages 1–1.
- Ramnath Balasubramanyan, William W. Cohen, Douglas Pierce, and David P. Redlawsk. 2012. Modeling polarizing topics: When do different political communities respond differently to the same news? In John G. Breslin, Nicole B. Ellison, James G. Shanahan, and Zeynep Tufekci, editors, Proceedings of the Sixth International Conference on Weblogs and Social Media, Dublin, Ireland, June 4-7, 2012. The AAAI Press.
- Jonah A. Berger and Katherine L. Milkman. 2012. What makes online content viral? Journal of Marketing Research, pages 192–205, 08.
- Margaret M. Bradley and Peter J. Lang. 1994. Measuring emotion: the self-assessment manikin and the semantic differential. Journal of behavior therapy and experimental psychiatry, 25 1:49–59.
- Xiao Chen, Changlong Sun, Jingjing Wang, Shoushan Li, Luo Si, Min Zhang, and Guodong Zhou. 2020. Aspect sentiment classification with document-level sentiment preference modeling. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 3667–3677, Online, July. Association for Computational Linguistics.
- Dorottya Demszky, Nikhil Garg, Rob Voigt, James Zou, Jesse Shapiro, Matthew Gentzkow, and Dan Jurafsky. 2019. Analyzing polarization in social media: Method and application to tweets on 21 mass shootings. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers), pages 2970–3005. Association for Computational Linguistics.
- Marco Guerini and Jacopo Staiano. 2015. Deep feelings: A massive cross-lingual study on the relation between emotions and virality. In Proceedings of the 24th International Conference on World Wide Web, WWW '15 Companion, page 299–305, New York, NY, USA. Association for Computing Machinery.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In Marilyn A. Walker, Heng Ji, and Amanda Stent, editors, Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers), pages 2227–2237. Association for Computational Linguistics.
- James Russell. 1980. A circumplex model of affect. Journal of Personality and Social Psychology, 39:1161–1178, 12.
- Carlo Strapparava and Rada Mihalcea. 2007. SemEval-2007 task 14: Affective text. In Proceedings of the Fourth International Workshop on Semantic Evaluations (SemEval-2007), pages 70–74, Prague, Czech Republic, June. Association for Computational Linguistics.
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 english lemmas. Behavior Research Methods, 45:1191–1207.
- Chen Zhang, Qiuchi Li, and Dawei Song. 2019. Aspect-based sentiment classification with aspect-specific graph convolutional networks. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 4568–4578, Hong Kong, China, November. Association for Computational Linguistics.