

What does it mean to be a Wutbürger? A first exploration.

Manfred Klenner

Institute of Computational Linguistics
Andreasstrasse 15, 8050 Zurich, Switzerland
klenner@cl.uzh.ch

Abstract. In this paper, we undertake an attempt to characterize the world view of what are called Wutbürger in Germany, that is citizen who are enraged by the current political and social situation. In order to find out what makes a Wutbürger a Wutbürger, we analyze Facebook posts on the basis of a lexical resource where nouns, adjectives and verbs are classified according to Plutchik's primary emotions. We also introduce new polar roles of verbs that help to identify the writer perspective. This way, we are able to identify targets and the Wutbürger's stance towards them. As textual data, we utilize about 100,000 Facebook posts of a German right-wing party whose members are obvious exemplars of the notion of a Wutbürger.

1 Introduction

In a number of western societies populism has (re)entered the scene and especially rampages in the social media. Hate speech, shit storms etc. are extreme forms of such undemocratic tendencies. In Germany, the notion of a *Wutbürger* has been coined, that is, citizen who are disappointed by the government and the social situation and their (verbal) behavior seem to be driven by rage (German *Wut*). A new, right-wing party evolved, the AfD (Alternative für Deutschland). We have access to about 100,000 Facebook posts of the AfD including reader comments that mostly stem from AfD proponents - who clearly form a subset of German Wutbürger. Our research question was: Can we find out, how a Wutbürger perceives the world and what, after all, is the objective of his Wut.

A first step towards this goal is to measure the emotional fingerprint of the texts produced by Wutbürger and compare it to the fingerprint of a related text genre. We use the Tübingen (German newspaper) Treebank (TüBa-D/Z) [11] as a reference corpus. In order to compare the emotional load of the AfD texts (303,563 sentences) and the TüBa-D/Z texts (95,595 sentences), we perform a lexicon-based analysis. That is, we count the primary emotions by the use of words that are indicative of these emotions. This is a straightforward approach, but it should tell us reliably what the prevalent emotions in these texts are and whether the AfD texts are more loaded than the newspaper reference texts.

The emotional fingerprint does not tell us anything about the world view of a Wutbürger: who are his proponents and who are his opponents? We would like

to exploit the idea that the identification of these targets can be supported and accomplished by a more fine-grained classification of lexical items. We not only assign primary emotions to words (verbs, nouns and adjectives), we also identify those words that have an implicit writer perspective and explicate which one it is. For instance, the adjective *ineffable* in a phrase like *the ineffable chancellor* expresses that the writer has a negative attitude towards the referent of the noun and a sentence like *Merkel jerks the German citizen around* allows to infer that the writer believes that the referent at subject position is an immoral actor, a cheater one might say. Also, the direct object is perceived as a victim of the cheater. Since the writer is against the cheater, he is in favor of the victim.

2 Lexical Resources

Starting with the freely available lexicons described in [2]¹ and [4]² we identified those verbs, nouns and adjectives that have an emotional dimension (e.g. *to love, to hate, gratitude, joy, pleasant, happy*). We then classified each of the 168 verbs, 225 nouns and 300 adjectives according to Plutchik’s [7] eight primary emotions which are *anger, fear, sadness, disgust, surprise, anticipation, trust,* and *joy*. This was done by two annotators, who achieved a Kappa value of κ 0.73. Furthermore, we annotated those verbs, nouns and adjectives that refer to moral, e.g. *to lie, donation*. See Figure 1 for an overview (*pos, neg* are shortcuts for *positive, negative* respectively). Kappa was $\kappa = 0.66$.

	pos emotion	neg emotion	pos moral	neg moral	pos factual	neg factual	#
verb	49	119	15	71	553	1170	1977
adj	118	182	286	569	1010	1103	3268
noun	91	134	104	436	663	1229	2657

Fig. 1. Lexicon Overview: Word Frequencies

The columns for *factual* denote words that are positive or negative without reference to either (a particular) emotion or moral. We could say that they are positive or negative on a factual level. For instance *to sicken, recover, congratulate* are examples of such verbs, whereas *mistake, disease, transparency, security, right, wrong* are examples for such nouns and adjectives. This is a crucial distinction: such words do not indicate a writer perspective, but the contribute to polarity decisions, nevertheless.

We took the 254 verbs classified as either belonging to the emotion or moral dimension as a basis for further annotations. We identified 58 verbs with a very strong writer perspective either on the actor or the experiencer role or on both (*to cheat, to jerk sb around*). We then coined for the six verb classes derived that way special role labels. The set of agent roles is: prole, baiter, hater, torturer, hypocrite, choleric. Experiencer roles are sufferer and victim. To give an example: the verb *flare up* (aufbrausen) bears the emotion *anger* and the semantic

¹ <http://bics.sentimental.li/files/8614/2462/8150/german.lex>

² https://pub.cl.uzh.ch/projects/opinion/lrec_data.txt

role of the subject is that of *choleric*. Our hypothesis is that these roles better capture the writer perspective, since they express how the writer conceptualizes these referents. Note that we assign these roles to subcategorization frames, not to verbs. We specified these verbs along the line proposed by [4]. That is, we modeled the various subcategorization frames of a verb and assigned it a polar effect (positive or negative) and for some of the verbs also a dedicated polar role (sufferer, torturer etc.). We thus were able to find out who the AfD believes to be a torturer, a baiter etc. and who suffers from the situation described.

3 Corpus Statistics and Lexical Coverage

Our present endeavor is basically one that exploits an existing, but carefully refined and augmented lexical resource. Especially our new verb classes with a new kind of polar roles are meant to make the writer perspective more nuanced. We are at the very beginning of a sophisticated study. At the moment, however, there is no gold standard and thus no machine learning involved.

We found 9,012 verb types in the Facebook posts, which gave us altogether 419,034 verb tokens in 303,563 sentences (word tokens altogether: 5,249,613). The 1052 verb types of our lexicon found in the data (61 verbs did not occur), amounts to 83,658 verb tokens, which is - taking into account that one sentence might have more than one model verb - about 25% coverage (a model verb in each 4th sentence). Our verb resource seems to have a good coverage, thus. If we just look at moral verbs, we get 11,153 hits, 64 of the 85 moral verb types do occur in the posts. The 168 emotional verbs occur with a frequency of 17,102.

In order to quantify the emotional load of the Facebook posts, we used the Tübinger Treebank (TüBa-D/Z) as a reference corpus. The TüBa-D/Z comprises 95,595 sentences. The coverage of our verb resource is again quite good: we found 930 verb types with 22,679 verb tokens, which is a coverage of 23.79% (again almost each 4th sentence bears a model verb).

4 Emotional Fingerprint

We use our emotion lexicon in order to diagnose the emotional load of the AfD posts. In Fig. 2, we compare the AfD posts and the newspaper text wrt. to the emotions present. We determined the frequency of words belonging to a particular emotion and normalized by the total number of emotion words (found in the posts) (left hand side), and by the number of sentences (right hand side), respectively.

As we can see from Fig. 2, fear is the most prominent emotion of a Wutbürger and not anger (a prestige of *rage*) while at the same time *sadness* is not a prevalent emotion of a Wutbürger. All other emotions are almost identically distributed in both, AfD posts and newspaper text. Our expectation, namely that the AfD posts would have a higher emotional load than news texts, was not confirmed.

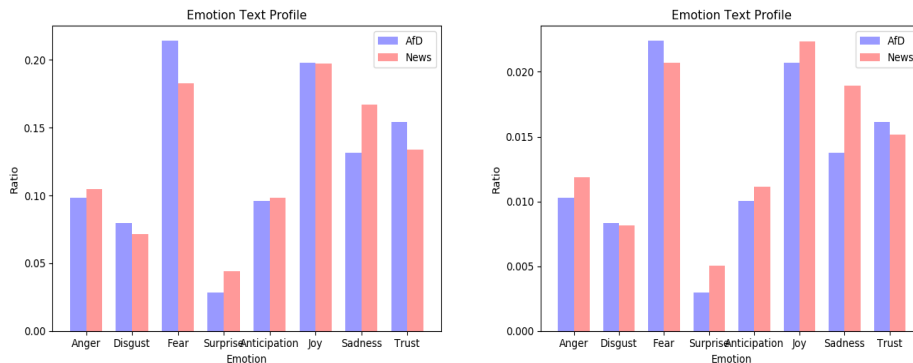


Fig. 2. Emotional Fingerprints of AfD Facebook Post and a Newspaper Corpus

We also had a look at the moral dimension. The TüBa-D/Z refers to 7490 nouns and adjectives classified as positive (32%) or negative (60%) from a moral perspective, i.e. 7.8% of the sentences refer to that dimension. In the AfD texts, 32167 tokens were found, which is about 10.6% (73% negative, 27% positive). This clearly shows that (negative) moral argumentation is a central attitude of a Wutbürger. If we have a look at verbs the picture is similar: 1.6% of the news texts contain a moralizing verb, whereas 2.3% of the AfD posts do so.

5 Target Identification and Stance Analysis

A *polar role* is the label for the logical subject (agent) or object (theme, patient, experiencer) that indicates the positive or negative role its filler plays. The inventory of polar roles is not fixed, yet. We have defined a couple of fine-grained polar roles that are meant to indicate a more nuanced writer perspective. These roles are *baiter*, *hater*, *choleric*, *hypocrite*, *prole*, *torturer* and *sufferer*, *victim*. The definition of these roles is straightforward: we just had to fix the corresponding verbs and determine which semantic role bears which polar role. Take the polar role *prole*. There is a number of verbs in German (we have identified 18) that indicate that the writer implicitly classifies the agent of such a verb as a prole (*anlabern* (to chat so up), *anpöbeln* (to accost sb)). Thus, the agents of such verbs are negative targets from the point of view of the writer. It turns out that in the AfD texts journalists, do-gooder, politicians, asylum seekers, the print media are, among others, conceptualized as proles.

In order to derive these writer perspectives, we have parsed the AfD posts with a dependency parser [10], normalized the parse trees (e.g. passive voice) and extracted the filler of the polar roles. This way we found e.g. that (the targets) Merkel, the German government, the police and the press are baiter, hater and torturers etc. The AfD, the German citizen and women were victims and sufferer. Clearly, these lists are not perfect. There are (third person) pronouns in it and also words denoting non-actors. The goal of this explorative study was to get a

proof of concept not a full-fledged evaluation. Nevertheless, we have carried out a small evaluation in order to find out where the noise comes from. We randomly took 50 sentences where the German chancellor Angela Merkel was the logical subject of a morally negative verb (e.g. *cheat, threaten, diss, violate*) and 50 where the AfD was the experiencer or patient of such a verb. Only 9 out of the 100 decisions were wrong due to 4 parsing errors and 5 modal constructions that erroneously passed our modal filter.

An interesting finding concerns the role of the AfD (i.e. Wutbürger) itself. If we look at those who are hated, we get: Arabs, strangers, Merkel, Muslims, comrades but also Germany and the AfD. A closer inspection reveals that the Wutbürger do not disguise or veil their rage. They use verbs with AfD (or *I* or *we*) as agents that indicate that they are haters.

Our verb resource also allows for more sophisticated inferences. We have coined the notion of a violator of morality for the following set of actors: the set of actors of a verb that casts a negative effect on its object which is - according to the polarity lexicon - positive:

$$\lambda X.(\exists Verb, Y: subj(Verb, X) \wedge effect(Verb, obj, neg) \wedge obj(Verb, Y) \wedge polarity(Y, pos))$$

An example is *Merkel destroys the security of Germany* where *security* is positive and *destroy* casts a negative effect on the direct object (obj)

If we, however, change *polarity(Y, pos)* to *polarity(Y, neg)* than we get a strong proponent of the AfD: to disapprove something negative is positive.

6 Related Work

One topic of this paper is lexicon-based, document-level emotion detection. For an overview of similar approaches see e.g. [1]. We have specified the first German emotion lexicon, where words are associated with primary emotions and - if applicable - writer perspectives.

The role verbs play in sentiment analysis and stance detection has received increased attention over the last years, cf. [6], [9], [3], [8], [5]. The main difference to our German verb resource is that we not only specify the polar effects (positive or negative) a verb casts on its semantic roles, we also strive to assign fine-grained role labels such as *torturer* etc. Again this is meant to allow for a finer nuanced writer perspective, which not only helps to identify targets, but also the stance taken towards those targets. Another distinctive feature is that we combine bottom-up and top-down information in order to derive stance (see last section).

7 Conclusions

We have introduced two new resources for German: a fine-grained verb resource with polar roles that reveal the writer perspective, and an emotion lexicon where words are classified as one of eight primary emotions. Also words related to moral are specified. We used this in order to find out whether Wutbürger texts do have a

clear emotional fingerprint compared to news texts. We found that *fear* (and not *Wut*, i.e. *rage*) is the prevalent emotion and that Wutbürger significantly more often argue on the basis of moral than a reference newspaper corpus. Another insight is that Wutbürger do not hide their rage (in their own sentences they often occupy negative polar roles such as hater). More sophisticated search pattern on the basis of top down and bottom up restrictions give rise to interesting inferences (someone who disapproves something positive is a violator of morality). All this is meant as a first explorative study: is our lexicon large enough to be useful, is our fine-grained verb resource broadly applicable. A fuller answer to the question raised in the title must await a thorough empirical investigation.

References

1. Haji Binali, Chen Wu, and Vidyasagar Potdar. Computational approaches for emotion detection from text. In *Proceedings of IEEE Intern. Conf. on Digital Ecosystems and Technologies*, 2010.
2. Simon Clematide and Manfred Klenner. Evaluation and extension of a polarity lexicon for German. In *Proceedings of the First Workshop on Computational Approaches to Subjectivity and Sentiment Analysis (WASSA)*, pages 7–13, 2010.
3. Lingjia Deng and Janyce Wiebe. Joint prediction for entity/event-level sentiment analysis using probabilistic soft logic models. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 179–189, 2015.
4. Manfred Klenner and Michael Amsler. Sentiframes: A resource for verb-centered German sentiment inference. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, Paris, France, may 2016. European Language Resources Association (ELRA).
5. Manfred Klenner, Don Tuggener, and Simon Clematide. Stance detection in Facebook posts of a German right-wing party. In *LSDSem 2017/LSD-Sem Linking Models of Lexical, Sentential and Discourse-level Semantics*, April 2017.
6. Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. Semantically distinct verb classes involved in sentiment analysis. In *IADIS AC (1)*, pages 27–35, 2009.
7. Robert Plutchik. *A general psychoevolutionary theory of emotion*. Academic press, NewYork, 1980.
8. Hannah Rashkin, Sameer Singh, and Yejin Choi. Connotation frames: A data-driven investigation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, Berlin, Germany, August 2016.
9. Kevin Reschke and Pranav Anand. Extracting contextual evaluativity. In *Proc. of the Ninth International Conf. on Computational Semantics*, pages 370–374, 2011.
10. Rico Sennrich, Martin Volk, and Gerold Schneider. Exploiting synergies between open resources for German dependency parsing, POS-tagging, and morphological analysis. In *Recent Advances in Natural Language Processing (RANLP 2013)*, pages 601–609, September 2013.
11. Heike Telljohann, Erhard W. Hinrichs, Sandra Kübler, Heike Zinsmeister, and Kathrin Beck. Stylebook for the Tübingen treebank of written German. Technical report, Universität Tübingen, Seminar für Sprachwissenschaft, 2009.