

Changes in European Solidarity Before and During COVID-19: Evidence from a Large Crowd- and Expert-Annotated Twitter Dataset

Alexandra Ils[†], Dan Liu[‡], Daniela Grunow[†], Steffen Eger[‡]

[†] Department of Social Sciences, Goethe University Frankfurt am Main

[‡] Natural Language Learning Group, Technische Universität Darmstadt

{ils, grunow}@soz.uni-frankfurt.de,

dan.liu.19@stud.tu-darmstadt.de, eger@aiphes.tu-darmstadt.de

Abstract

We introduce the well-established social scientific concept of social solidarity and its contestation, anti-solidarity, as a new problem setting to supervised machine learning in NLP to assess how European solidarity discourses changed before and after the COVID-19 outbreak was declared a global pandemic. To this end, we annotate 2.3k English and German tweets for (anti-)solidarity expressions, utilizing multiple human annotators and two annotation approaches (experts vs. crowds). We use these annotations to train a BERT model with multiple data augmentation strategies. Our augmented BERT model that combines both expert and crowd annotations outperforms the baseline BERT classifier trained with expert annotations only by over 25 points, from 58% macro-F1 to almost 85%. We use this high-quality model to automatically label over 270k tweets between September 2019 and December 2020. We then assess the automatically labeled data for how statements related to European (anti-)solidarity discourses developed over time and in relation to one another, before and during the COVID-19 crisis. Our results show that solidarity became increasingly salient and contested during the crisis. While the number of solidarity tweets remained on a higher level and dominated the discourse in the scrutinized time frame, anti-solidarity tweets initially spiked, then decreased to (almost) pre-COVID-19 values before rising to a stable higher level until the end of 2020.

1 Introduction

Social solidarity statements and other forms of collective pro-social behavior expressed in online media have been argued to affect public opinion and political mobilization (Fenton, 2008; Margolin and Liao, 2018; Santhanam et al., 2019; Tufekci, 2014). The ubiquity of social media enables individuals to feel and relate to real-world problems through

solidarity statements expressed online and to act accordingly (Fenton, 2008). Social solidarity is a key feature that keeps modern societies integrated, functioning and cohesive. It constitutes a moral and normative bond between individuals and society, affecting people’s willingness to help others and share own resources beyond immediate rational individually-, group- or class-based interests (Silver, 1994). National and international crises intensify the need for social solidarity, as crises diminish the resources available, raise demand for new and additional resources, and/or require readjustment of established collective redistributive patterns, e.g. inclusion of new groups. Because principles of inclusion and redistribution are contested in modern societies and related opinions fragmented (Fenton, 2008; Sunstein, 2018), collective expressions of social solidarity online are likely contested. Such statements, which we refer to as anti-solidarity, question calls for social solidarity and its framing, i.e. towards whom individuals should show solidarity, and in what ways (Wallaschek, 2019).

For a long time, social solidarity was considered to be confined to local, national or cultural groups. The concept of a European society and European solidarity (Gerhards et al., 2019), a form of solidarity that goes beyond the nation state, is rather new. European solidarity gained relevance with the rise and expansion of the European Union (EU) and its legislative and administrative power vis-à-vis the EU member states since the 1950s (Baglioni et al., 2019; Gerhards et al., 2019; Koos and Seibel, 2019; Lahusen and Grasso, 2018). After decades of increasing European integration and institutionalization, the EU entered into a continued succession of deep crises, beginning with the European Financial Crisis in 2010 (Gerhards et al., 2019). Experiences of recurring European crises raise concerns regarding the future of European society and its foundation, European solidarity. Eurosceptics

and right-wing populists claim that social solidarity is, and should be, confined within the nation state, whereas supporters of the European project see European solidarity as a means to overcome the great challenges imposed on EU countries and its citizens today (Gerhards et al., 2019). To date, it is an open empirical question how strong and contested social solidarity really is in Europe, and how it has changed since the onset of the COVID-19 pandemic. Against this background, we ask whether we can detect changes in the debates on European solidarity before and after the outbreak of COVID-19. Our contributions are:

- (i) We provide a novel Twitter corpus annotated for expressions of social solidarity and anti-solidarity. Our corpus contains 2.3k human-labeled tweets from two annotation strategies (experts vs. crowds). Moreover, we provide over 270k automatically labeled tweets based on an ensemble of BERT classifiers trained on the expert and crowd annotations.
- (ii) We train BERT on crowd- and expert annotations using multiple data augmentation and transfer learning approaches, achieving over 25 points improvement over BERT trained on expert annotations alone.
- (iii) We present novel empirical evidence regarding changes in European solidarity debates before and after the outbreak of the COVID-19 pandemic. Our findings show that both expressed solidarity and anti-solidarity escalated with the occurrence of incisive political events, such as the onset of the first European lockdowns.

Our data and code are available from <https://github.com/lalashiwoya/socialsolidarityCOVID19>.

2 Related work

Social Solidarity in the Social Sciences. In the social sciences, social solidarity has always been a key topic of intellectual thought and empirical investigation, dating back to seminal thinkers such as Rousseau and Durkheim (Silver, 1994). Whereas earlier empirical research was mostly confined to survey-based (Baglioni et al., 2019; Gerhards et al., 2019; Koos and Seibel, 2019; Lahusen and Grasso, 2018) or qualitative approaches (Franceschelli, 2019; Gómez Garrido et al., 2018; Heimann et al., 2019), computational social science just started tackling concepts as complex as solidarity as part

of natural language processing (NLP) approaches (Santhanam et al., 2019).

In (computational) social science, several studies investigated the European Migration Crisis and/or the Financial Crisis as displayed in media discourses. These studies focused on differences in perspectives and narratives between mainstream media and Twitter, using topic models (Nerghees and Lee, 2019), and the coverage and kinds of solidarity addressed in leftist and conservative newspaper media (Wallaschek, 2019, 2020a), as well as relevant actors in discourses on solidarity, using discourse network measures (Wallaschek, 2020b). While these studies offer insight into solidarity discourses during crises, they all share a strong focus on mainstream media, which is unlikely to publicly reject solidarity claims (Wallaschek, 2019). Social media, in contrast, allows its users to perpetuate, challenge and open new perspectives on mainstream narratives (Nerghees and Lee, 2019). A first attempt to study solidarity expressed by social media users during crises has been presented by Santhanam et al. (2019). They assessed how emojis are used in tweets expressing solidarity relating to two crises through hashtag-based manual annotation—ignoring actual content of the tweets—and utilizing a LSTM network for automatic classification. Their approach, while insightful, provides a rather simple operationalization of solidarity, which neglects its contested, consequential and obligatory aspects vis-à-vis other social groups.

The current state of social science research on European social solidarity poses a puzzle. On the one hand, most survey research paints a rather optimistic view regarding social solidarity in the EU, despite marked cross-national variation (Binner and Scherschel, 2019; Dragolov et al., 2016; Gerhards et al., 2019; Lahusen and Grasso, 2018). On the other hand, the rise of political polarization and Eurosceptic political parties (Baker et al., 2020; Nicoli, 2017) suggests that the opinions, orientations and fears of a potentially growing political minority is underrepresented in this research. People holding extreme opinions have been found to be reluctant to participate in surveys and adopt their survey-responses to social norms (social desirability bias) (Bazo Vienrich and Creighton, 2017; Heerwegh, 2009; Janus, 2010). Research indicates that such minorities may grow in times of crises, with both short-term and long-term effects for public opinion and political trust (Gangl and Giustozzi,

2018; Nicoli, 2017). Our paper addresses these problems by drawing on large volumes of longitudinal social media data that reflect potential fragmentation of political opinion (Sunstein, 2018) and its change over time. Our approach will thus uncover how contested European solidarity is and how it developed since the onset of COVID-19.

Emotion and Sentiment Classification in NLP.

In NLP, annotating and classifying text (in social media) for sentiment or emotions is a well-established task (Demszky et al., 2020; Ding et al., 2020; Haider et al., 2020; Hutto and Gilbert, 2014; Oberländer and Klinger, 2018). Importantly, our approach focuses on expressions of (anti-)solidarity: For example, texts containing a positive sentiment towards persons, groups or organizations which are at their core anti-European, nationalistic and excluding reflect anti-solidarity and are annotated as such. Our annotations therefore go beyond superficial assessment of sentiment. In fact, the correlation between sentiment labels—e.g., as obtained from Vader (Hutto and Gilbert, 2014)—and our annotations in §3 is only ~ 0.2 . Specifically, many tweets labeled as solidarity use negatively connoted emotion words.

3 Data and Annotations

We use the unforeseen onset of the COVID-19 crisis, beginning with the first European lockdown, enacted late February to early March 2020, to analyze and compare social solidarity data before and during the COVID-19 crisis as if it were a natural experiment (Creighton et al., 2015; Kuntz et al., 2017). In order to utilize this strategy and keep the baseline solidarity debate comparable before and after the onset of the COVID-19 crisis, we confined our sample to tweets with hashtags predominantly relating to two previous European crises whose effects continue to concern Europe, its member states and citizens: (i) Migration and the distribution of refugees among European member states, and (ii) Financial solidarity, i.e. financial support for indebted EU countries. The former solidarity debate predominantly refers to the Refugee Crisis since 2015 and the living situation of migrants, the latter mostly relates to the Financial Crisis, followed by the Euro Crisis, and concerns the excessive indebtedness of some EU countries since 2010.¹

¹Further analyses (not shown) revealed that around 20 percent of the tweets in our sample relate to solidarity regarding other issues.

Data. We crawled 271,930 tweets between 01.09.2019 and 31.12.2020, written in English or German and geographically restricted to Europe, to obtain setups comparable to the survey-based social science literature on European solidarity. We only crawled tweets that contained specific hashtags, to filter for our two topics, i.e. refugee and financial solidarity. We started with an initial list of hashtags (e.g., “#refugeecrisis”, “#eurobonds”), which we then expanded via co-occurrence statistics. We manually evaluated 456 co-occurring hashtags with at least 100 occurrences to see if they represented the topics we are interested in. Ultimately, we selected 45 hashtags (see appendix) to capture a wide range of the discourse on migration and financial solidarity. Importantly, we keep the hashtag list associated with our 270k tweets constant over time.²

Definition of Social Solidarity. In line with social scientific concepts of social solidarity, we define social solidarity as expressed and/or called for in online media as “the preparedness to share one’s own resources with others, be that directly by donating money or time in support of others or indirectly by supporting the state to reallocate and redistribute some of the funds gathered through taxes or contributions” (Lahusen and Grasso, 2018, p. 4). We define anti-solidarity as expressions that contest this type of social solidarity and/or deny solidarity towards vulnerable social groups and other European states, e.g. by promoting nationalism or the closure of national borders (Burgoon and Rooduijn, 2021; Cinalli et al., 2020; Finseraas, 2008; Walaschek, 2017).

Expert Annotations. After crawling and preparing the data, we set up guidelines for annotating tweets. Overall, we set four categories to annotate, with solidarity and anti-solidarity being the most important ones. A tweet indicating support for people in need, the willingness and/or gratitude towards others to share resources and/or help them is considered expressing *solidarity*. The same applies to tweets criticizing the EU in terms of not doing enough to share resources and/or help socially vulnerable groups as well as advocating for the EU as a solidarity union. A tweet is considered to be expressing *anti-solidarity* statements

²We follow a purposeful sampling frame, but this necessarily introduces a bias in our data. While we took care of including a variety of hashtags, we do not claim to have captured the full extent of discourse concerning the topics migration and financial solidarity.

if the above-mentioned criteria are reversed, and/or, the tweet contains tendencies of nationalism or advocates for closed borders. Not all tweets fit into these classes, thus we introduce two additional categories: `ambivalent` and `not applicable`. While the `ambivalent` category refers to tweets that could be interpreted as both expressing solidarity and anti-solidarity statements, the second category is reserved for tweets that do not contain the topic of (anti-)solidarity at all or refer to topics that are not concerned with discourses on refugee or financial solidarity. Table 1 contains example tweets for all categories. Full guidelines for the annotation of tweets are given in the appendix.

We divided the annotation process into six working stages (I-VI) to refine our data set and annotation standards over time and strengthen inter-annotator reliability through subsequent discussions among annotators and social science experts. Our annotators included four university students majoring in computer science, one computer science faculty member as well as two social science experts (one PhD student and one professor). We started the training of seven annotators with a small dataset that they annotated independently and refined the guidelines during the annotation process. In the training period, which lasted three iterations (I-III), we achieved Cohen’s kappa values of 0.51 among seven annotators. In working stage IV, two groups of two annotators annotated 339 tweets with hashtags not included before. Across the four annotators, Cohen’s kappa values of 0.49 were reached. In working stages V and VI, one group of two students annotated overall 588 tweets, with a resulting kappa value of 0.79 and 0.77 respectively.

While the kappa value was low in the first stages, we managed to raise the inter-annotator reliability over time through discussions with the social science experts and extension of the guidelines. We also introduced a gold-standard for annotations from stage II onward which served as orientation. This was determined by majority voting and discussions among the annotators. For cases where a decision on the gold-standard label could not be reached, a social science expert decided on the gold-standard label; some hard cases were left undecided (not included in the dataset).

The gold-standard additionally served as human reference performance which we compared the model against. On average across all stages, our kappa agreement is 0.64 for four and 0.69 for

three classes (collapsing `ambivalent` and `not applicable`), while the macro F1-score is 69% for four and 78.5% for three classes. However, in the final stages, the agreement is considerably higher: above 80% macro-F1 for four and between 85.4% and 89.7% macro-F1 for three classes.

Crowd annotations. We also conducted a ‘crowd experiment’ with students in an introductory course to NLP. We provided students with the guidelines and 100 expert annotated tweets as illustrations. We trained crowd annotators in three iterations. 1) They were assigned reading the guidelines and looking at 30 random expert annotations. Then they were asked to annotate 20 tweets themselves and self-report their kappa agreement with the experts (we provided the labels separately so that they could further use the 20 tweets to understand the annotation task). 2) We repeated this with another 30 tweets for annotator training and 20 tweets for annotator testing. 3) They received 30 expert-annotated tweets for which we did not give them access to expert labels, and 30 entirely novel tweets, that had not been annotated before. These 60 final tweets were presented in random order to each student. 50% of the 30 novel tweets were taken from before September 2020 and the other 50% were taken from after September 2020.

125 students participated in the annotation task. The annotation experiment was part of a bonus the students could achieve for the course (counted 12.5% of the overall bonus for the class). Each novel tweet was annotated by up to 3 students (2.7 on average). To obtain a unique label for each crowd-annotated tweet, we used the following simple strategy: we either chose the majority label among the three annotators or the annotation of the most reliable annotator in case there was no unique majority label. The annotator that had the highest agreement with the expert annotators was taken as most reliable annotator.




Figure 1: Distribution of kappa agreements of crowd workers with expert annotated gold-standard, 3 classes.

Children caught up in the Moria camp fire face unimaginable horrors #SafePassage #RefugeesWelcome	solidarity
Most people supporting #RefugeesWelcome are racists or psychopaths	anti-solidarity
Does this rule apply to every UK citizen as well as every #AsylumSeeker?	ambivalent
Let’s make #VaccinesWork for everyone #LeaveNoOneBehind	not applicable

Table 1: Paraphrased (and translated) sample of annotated tweets in our dataset, together with labels.

	S	A	AMB	NA	Total
Experts	386	246	113	174	919
Crowds	768	209	186	217	1380

Table 2: Number of annotated tweets (after geofiltering) for the four classes `solidarity` (S), `anti-solidarity` (A), `ambivalent` (AMB), and `not applicable` (NA).

Kappa agreements of students with the experts are shown in Figure 1. The majority of students has a kappa agreement with the gold-standard of between 0.6-0.7 when three classes are taken into account and between 0.5-0.6 for four classes.

In Table 2, we further show statistics on our annotated datasets: we have 2299 annotated tweets in total, about 60% of which have been annotated by crowd-workers. About 50% of all tweets are annotated as `solidarity`, 20% as `anti-solidarity`, and 30% as either `not-applicable` or `ambivalent`. In our annotations, 1196 tweets are English and 1103 are German.³ Finally, we note that the distribution of labels for expert and crowd annotations are different, i.e., the crowd annotations cover more `solidarity` tweets. The reason is twofold: (a) for the experts, we oversampled hashtags that we believed to be associated more often with `anti-solidarity` tweets as the initial annotations indicated that these would be in the minority, which we feared to be problematic for the automatic classifiers. (b) The time periods in which the tweets for the experts and crowd annotators fall differ.

4 Methods

We use multilingual BERT (Devlin et al., 2019) / XLM-R (Conneau et al., 2020) to

³In our automatically labeled data, the majority of tweets is German. We assumed all German tweets to come from within the EU, while the English tweets would be geofiltered more aggressively.

classify our tweets in a 3-way classification problem (`solidarity`, `anti-solidarity`, `other`), not differentiating between the classes `ambivalent` and `non-applicable` since our main focus is on the analysis of changes in (anti-)solidarity. We use the baseline MBERT model: `bert-base-multilingual-cased` and the base XLM-R model: `xlm-roberta-base`. We implemented several data augmentation/transfer learning techniques to improve model performance:

- **Oversampling of minority classes:** We randomly duplicate (expert and crowd annotated) tweets from minority classes until all classes have the same number of tweets as the majority class `solidarity`.
- **Back-translation:** We use the Google Translate API to translate English tweets into a pivot language (we used German), and pivot language tweets back into English (for expert and crowd-annotated tweets).
- **Fine-tuning:** We fine-tune MBERT / XLM-R with masked language model and next sentence prediction tasks on domain-specific data, i.e., our crawled unlabeled tweets.
- **Auto-labeled data:** As a form of self-learning, we train 9 different models (including oversampling, back-translation, etc.) on the expert and crowd-annotated data, then apply them to our full dataset (of 270k tweets, see below). We only retain tweets where 7 of 9 models agree and select 35k such tweets for each label (`solidarity`, `anti-solidarity`, `other`) into an augmented training set, thus increasing training data by 105k auto-labeled tweets.
- **Ensembling:** We take the majority vote of 15 different models to leverage heterogeneous information. The $k = 15$ models, like the $k = 9$ models above, were determined as the top- k models by their dev set performance.

We also experimented with re-mapping multilin-

gual BERT and XLM-R (Cao et al., 2020; Zhao et al., 2020a,b) as they have not seen parallel data during training, but found only minor effects in initial experiments.

5 Experiments

In §5.1, we describe our experimental setup. In §5.2, we show the classification results of our baseline models on the annotated data and the effects of our various data augmentation and transfer learning strategies. In §5.3, we analyze performances of our best-performing models. In §5.4, we automatically label our whole dataset of 270k tweets and analyze changes in solidarity over time.

5.1 Experimental Setup

To examine the effects of various factors, we design several experimental conditions. These involve (i) using only hashtags for classification, ignoring the actual tweet text, (ii) using only text, without the hashtags, (iii) combining expert and crowd annotations for training, (iv) examining the augmentation and transfer learning strategies, (v) ensembling various models using majority voting.

All models are evaluated on randomly sampled test and dev sets of size 170 each. Both dev and test set are taken from the expert annotations. We use the dev set for early stopping. To make sure our results are not an artefact of unlucky choices of test and dev sets, we report averages of 3 random splits where test and dev set contain 170 instances in each case (for reasons of computational costs, we do so only for selected experimental conditions).

We report the macro-F1 score to evaluate the performance of different models. Hyperparameters of our models can be found in our github.

5.2 Results

The main results are reported in Table 3. Using only hashtags and expert annotated data yields a macro-F1 score of below 50% for MBERT and XLM-R. Including the full texts improves this by over 8 points (almost 20 points for XLM-R). Adding crowd-annotations yields another substantial boost of more than 6 points for MBERT. Removing hashtags in this situation decreases the performance between 5 and 6 points. This means that the hashtags indeed contain import information, but the texts are more important than the hashtags: with hashtags only, we observe macro-F1 scores between 42 and 49%, whereas with text only the performances are

substantially higher, between 58 and 60%. While using hashtags only means less data since not all of our tweets have hashtags, the performance with only hashtags on the test sets stays below 50%, both with 572 and more than 1500 tweets for training.

Next, we analyze the data augmentation and transfer learning techniques. Including auto-labeled data drastically increases the train set, from below 2k instances to over 100k. Even though these instances are self-labeled, performance increases by over 13 points to about 78% macro-F1. Additionally oversampling or backtranslating the data does not yield further benefits, but pretraining on unlabeled tweets is effective even here and boosts performance to over 78%. Combining all strategies yields scores of up to almost 80%. Finally, when we consider our ensemble of 15 models, we achieve a best performance of 84.5% macro-F1 on the test set, close to the human macro-F1 agreement for the experts in the last rounds of annotation.

To sum up, we note: (i) adding crowd annotated data clearly helps, despite the crowd annotated data having a different label distribution; (ii) including text is important for classification as the classification with hashtags only performs considerably worse; (iii) data augmentation (especially self-labeling), combining models and transfer learning strategies has a further clearly positive effect.

5.3 Model Analysis

Our most accurate ensemble models perform best for the majority class *solidarity* with an F1-score of almost 90%, about 10 points better than for *anti-solidarity* and over 5 points better than for the *other* class. A confusion matrix for this best performing model is shown in Table 4. Here, *anti-solidarity* is disproportionately misclassified as either *solidarity* or the *other* class.

Table 5 shows selected misclassifications for our ensemble model with performance of about 84.5% macro-F1. This reveals that the models sometimes leverage superficial lexical cues (e.g., the German political party ‘AfD’ is typically associated with anti-solidarity towards EU and refugees), including hashtags (‘Remigration’); see Figure 2, where we used LIME (Ribeiro et al., 2016) to highlight words the model pays attention to. To further gain insight into the misclassifications, we had one social science expert reannotate all misclassifications. From

Condition	Train size	MBERT		XLM-R	
		Dev	Test	Dev	Test
E, Hashtag only	572	51.7±0.5	49.0±1.1	48.0±0.9	44.0±0.8
E	579	64.2±1.2	57.7±0.4	64.0	63.3
E+C	1959	66.4±0.5	64.0±1.5	65.0	64.8
E+C, No Hashtags	1959	64.0±0.3	58.0±0.5	62.0	60.0
E+C, Hashtag Only	1567	55.8±2.0	49.5±2.1	47.8	42.2
E+C+Auto label	106959	76.4	78.3	77.5	78.4
E+C+Auto label+Oversample	108048	76.4	76.3	77.4	76.9
E+C+Auto label+Backtranslation	108918	76.0	77.1	77.5	78.7
E+C+Auto label+Pretraining	106959	78.4	78.8	78.6	79.0
E+C+ALL	110007	78.8±1.3	78.6±0.8	78.9	79.7

Table 3: Macro-F1 scores (in %) for different conditions. Entries with \pm give averages and standard deviations over 3 different runs with different test and dev sets. ‘E’ stands for Experts, ‘C’ for Crowds. ‘ALL’ refers to all data augmentation and transfer learning techniques.

the 25 errors that our best model makes in the test set of 170 instances, the expert thinks that 12 times the gold standard is correct, 7 times the model prediction is correct, and in further 6 cases neither the model nor the gold standard are correct. This hints at some level of errors in our annotated data; it further supports the conclusion that our model is close to the human upper bound.

		Predicted		
		S	A	O
Gold	S	63	3	2
	A	5	37	4
	O	5	6	45

Table 4: Confusion matrix for best ensemble with macro-F1 score of 84.5% on the test set (for one specific train, dev, test split).

5.4 Temporal Analysis

Throughout the period observed in our data, discourses relating to migration were much more frequent than financial solidarity discourses. We crawled an average of 2526 tweets per week relating to migration (anti-)solidarity and an average of 174 financial (anti-)solidarity tweets, judging from the associated hashtags.

We used our best performing model to automatically label all our 270k tweets between September 2019 and December 2020. Solidarity tweets were about twice as frequent compared to anti-solidarity tweets, reflecting a polarized discourse in which solidarity statements clearly dominated. Figure

3 shows the frequency curves for solidarity, anti-solidarity and other tweets over time in our sample. The figure also gives the ratio

$$S/A := \frac{\# \text{Solidarity tweets}}{\# \text{Anti-Solidarity tweets}}$$

that shows the frequency of solidarity tweets relative to anti-solidarity tweets. Values above one indicate that more solidarity than anti-solidarity statements were tweeted that day.

Figure 3 displays several short-term increases in solidarity statements in our window of observation. Further analysis shows that these peaks have been immediate responses to drastic politically relevant events in Europe, which were also prominently covered by mainstream media, i.e. COVID-19-related news, natural disasters, fires, major policy changes. We illustrate this in the following.

On March 11th 2020, the World Health Organization (WHO) declared the COVID-19 outbreak a global pandemic. Shortly before and after, European countries started to take a variety of countermeasures, including stay-at-home orders for the general population, private gathering restrictions, and the closure of educational and childcare institutions (ECDC, 2020a). With the onset of these interventions, both solidarity and anti-solidarity statements relating to refugees and financial solidarity increased dramatically. At its peak at the beginning of March, anti-solidarity statements markedly outnumbered solidarity statements (we recorded 2189 solidarity tweets vs. 2569 anti-solidarity tweets on march 3rd). In fact, the period in early March 2020 is the only extended period in our data where anti-




Figure 2: Our best-performing model (macro-F1 of 84.5%) predicts *anti-solidarity* for the current example because of the hashtag #Remigration (according to LIME). The tweet, also given as translation in Table 5 (2) below, is overall classified as *other* in the gold standard, as it may be considered as expressing no determinate stance. Here, we hide identity revealing information in the tweet, but our classifier sees it.

Text	Gold	Pred.
(1) You can drink a toast with the AFD misanthropists #seenotrettung #NieMehrCDU	S	A
(2) Why is an open discussion about #Remigration (not) yet possible?	O	A
(3) Raped and Beaten, Lesbian #AsylumSeeker Faces #Deportation	A	O

Table 5: Selected misclassifications of best performing ensemble model. We consider the bottom tweet misclassified in the expert annotated data (correct would be *solidarity*). Tweets are paraphrased and/or translated.

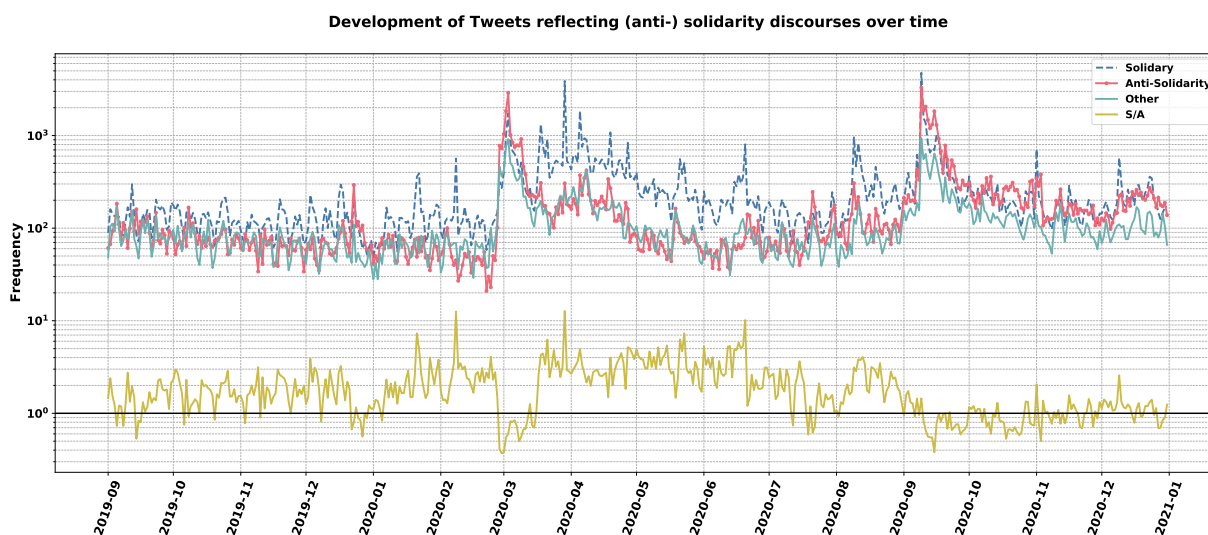


Figure 3: Frequency of *solidarity* (S), *anti-solidarity* (A) and *other* (O) tweets over time as well as the ratio S/A . The constant line 1 indicates where $S = A$. Y-axis is in log-scale.

solidarity statements outweighed *solidarity* statements. The dominance of *solidarity* statements was reestablished after two weeks. Over the following months, *anti-solidarity* statements decreased again to pre-COVID-19 levels, whereas *solidarity* statements remained comparatively high, with several peaks between March and September 2020.

Solidarity and *anti-solidarity* statements shot up again early-September 2020, with an unprecedented climax on September 9th. Introspection of our data shows that the trigger for this was the precarious situation of refugees after a fire destroyed the Mória Refugee Camp on the Greek island of Lesbos on the night of September 8th. Human Rights Watch had compared the camp to an

open-air prison in which refugees lived under inhumane conditions, and the disaster spurred debates about the responsibilities of EU countries towards refugees and the countries hosting refugee hot spots (i.e. Greece and Italy). At that time, COVID-19 infection rates in the EU were increasing but still low, and national measures to prevent the spread of infections relaxed in some and tightened in other EU countries (ECDC, 2020a,b). Further analyses (not displayed) show that the dominance of *solidarity* over *anti-solidarity* statements at the time was driven by tweets using hashtags relating to migration. The contemporaneous discourse on financial *solidarity* between EU countries was much less pronounced. From September 2020 to Decem-

ber 2020, solidarity and (anti-)solidarity statements were about equal in frequency, which means that anti-solidarity was on average on a higher level compared to the earlier time points in our time frame. This period also corresponds to the highest COVID-19 infection rates witnessed in the EU, on average, during the year 2020. In fact, the Spearman correlation between the number of anti-solidarity tweets in our data and infection rates is 0.45 and 0.47, respectively (infection rates within Germany and the EU); see Figure 4 in the appendix. Correlation with the number of solidarity tweets is, in contrast, non-significant.

Discussion Late February to mid-March 2020, EU governments began enacting lockdowns and other measures to contain COVID-19 infection rates, turning people’s everyday lives upside down. During this time frame, anti-solidarity statements peaked in our data, but solidarity statements quickly dominated thereafter again. During the summer of 2020, anti-solidarity tweets decreased whereas solidarity tweets continued to prevail on higher levels than before. A major peak on September 9th, in the aftermath of the destruction of the Mória Refugee Camp, signifies an intensification of the polarized solidarity discourse. From September to December 2020, anti-solidarity and solidarity statements were almost equal in number. Thus, the onset of the COVID-19 crisis as well as times of high infection rates concurred with disproportionately high levels of anti-solidarity, despite a dominance of solidarity overall. Whether the relationship between anti-solidarity and intensified strains during crises is indeed causal will be the scope of our future research.⁴

6 Conclusion

In this paper, we contributed the first large-scale human and automatically annotated dataset labeled for solidarity and its contestation, anti-solidarity. The dataset uses the textual material in social media posts to determine whether a post shows (anti-)solidarity with respect to relevant target groups. Our annotations, conducted by both trained experts and student crowd-workers, show overall good agreement levels for a challenging novel NLP task. We further trained augmented BERT models whose

⁴We made sure that the substantial findings reported here are not driven by inherently German (anti-)solidarity discourses. Still, our results are bound to the opinions of people posting tweets in the English and German language.

performance is close to the agreement levels of the experts and which we used for large-scale trend analysis of over 270k media posts before and after the onset of the COVID-19 pandemic. Our findings show that (anti-)solidarity statements climaxed momentarily with the first lockdown, but the predominance of solidarity expressions was quickly restored at higher levels than before. Solidarity and anti-solidarity statements were balanced by the end of the year 2020, when infection rates were rising.

The COVID-19 pandemic constitutes a worldwide crisis, with profound economic and social consequences for contemporary societies. It manifests yet another challenge for European solidarity, by putting a severe strain on available resources, i.e. national economies, health systems, and individual freedom. While the EU, its member countries and residents continued to struggle with the consequences of the Financial Crisis and its aftermath, as well as migration, the COVID-19 pandemic has accelerated the problems related to these former crises. Our data suggests that the COVID-19 pandemic has not severely negatively impacted the willingness of European Twitter users to take responsibility for refugees, while financial solidarity with other EU countries remained low on the agenda. Over time, however, this form of expressed solidarity became more controversial. On one hand, these findings are in line with survey-based, quantitative research and its rather optimistic overall picture regarding social solidarity in the EU during earlier crises (Baglioni et al., 2019; Gerhards et al., 2019; Koos and Seibel, 2019; Lahusen and Grasso, 2018); on the other hand, results from our correlation analysis suggests that severe strains during crises coincide with increased levels of anti-solidarity statements. We conclude that a convergence of opinion (Santhanam et al., 2019) among the European Twitter-using public regarding the target audiences of solidarity, and the limits of European solidarity vs. national interests, is not in sight. Instead, our widened analytic focus has allowed us to examine pro-social online behavior during crises and its opposition, revealing that European Twitter users remain divided on issues of European solidarity.

Acknowledgments

We thank the anonymous reviewers whose comments greatly improved the final version of the paper.

Ethical considerations. We will release only tweet IDs in our final dataset. The presented tweets in our paper were paraphrased and/or translated and therefore cannot be traced back to the users. No user identities of any annotator (neither expert nor crowd worker) will ever be revealed or can be inferred from the dataset. Crowd workers were made aware that the annotations are going to be used in further downstream applications and they were free to choose to submit their annotations. While our trained model could potentially be misused, we do not foresee greater risks than with established NLP applications such as sentiment or emotion classification.

References

- Simone Baglioni, Olga Biosca, and Tom Montgomery. 2019. [Brexit, division, and individual solidarity: What future for europe? Evidence from eight european countries.](#) *American Behavioral Scientist*, 63(4):538–550.
- Scott R Baker, Aniket Baksy, Nicholas Bloom, Steven J Davis, and Jonathan A Rodden. 2020. [Elections, political polarization, and economic uncertainty.](#) Working Paper 27961, National Bureau of Economic Research.
- Alessandra Bazo Vienrich and Mathew J. Creighton. 2017. [What’s left unsaid? In-group solidarity and ethnic and racial differences in opposition to immigration in the united states.](#) *Journal of Ethnic and Migration Studies*, 44(13):2240–2255.
- Kristina Binner and Karin Scherschel, editors. 2019. *Fluchtmigration und Gesellschaft: Von Nutzenkalkülen, Solidarität und Exklusion.* Arbeitsgesellschaft im Wandel. Beltz Juventa.
- Brian Burgoon and Matthijs Rooduijn. 2021. [‘Immigrationization’ of welfare politics? Anti-immigration and welfare attitudes in context.](#) *West European Politics*, 44(2):177–203.
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. [Multilingual alignment of contextual word representations.](#) *ICLR*.
- Manlio Cinalli, Olga Eisele, Verena K. Brändle, and Hans-Jörg Trenz. 2020. [Solidarity contestation in the public domain during the ‘refugee crisis’.](#) In Christian Lahusen, editor, *Citizens’ Solidarity in Europe*, pages 120–148.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Mathew J. Creighton, Amaney Jamal, and Natalia C. Malancu. 2015. [Has opposition to immigration increased in the united states after the economic crisis? an experimental approach.](#) *International Migration Review*, 49(3):727–756.
- Dorottya Demszky, Dana Movshovitz-Attias, Jeongwoo Ko, Alan Cowen, Gaurav Nemade, and Sujith Ravi. 2020. [GoEmotions: A dataset of fine-grained emotions.](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4040–4054, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding.](#) In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Keyang Ding, Jing Li, and Yuji Zhang. 2020. [Hash-tags, emotions, and comments: A large-scale dataset to understand fine-grained social emotions to online topics.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1376–1382, Online. Association for Computational Linguistics.
- Georgi Dragolov, Zsófia S. Ignác, Jan Lorenz, Jan Delhey, Klaus Boehnke, and Kai Unzicker. 2016. *Social cohesion in the western world: What holds societies together: Insights from the social cohesion radar.* Springer.
- ECDC. 2020a. [European centre for disease prevention and control, data on country response measures to covid-19.](#)
- ECDC. 2020b. [European centre for disease prevention and control, download historical data \(to 14 december 2020\) on the daily number of new reported covid-19 cases and deaths worldwide.](#)
- Natalie Fenton. 2008. [Mediating solidarity.](#) *Global Media and Communication*, 4(1):37–57.
- Henning Finseraas. 2008. [Immigration and preferences for redistribution: An empirical analysis of european survey data.](#) *Comparative European Politics*, 6(4):407–431.
- Michela Franceschelli. 2019. [Global migration, local communities and the absent state: Resentment and resignation on the italian island of lampedusa.](#) *Sociology*, 54(3):591–608.

- Markus Gangl and Carlotta Giustozzi. 2018. [The erosion of political trust in the great recession](#). *CORRODE Working Paper*, (5).
- Jürgen Gerhards, Holger Lengfeld, Zsófia S. Ignácz, Florian K. Kley, and Maximilian Priem. 2019. [European solidarity in times of crisis: Insights from a thirteen-country survey](#). Routledge.
- María Gómez Garrido, M. Antònia Carbonero Gamundí, and Anahí Viladrich. 2018. [The role of grassroots food banks in building political solidarity with vulnerable people](#). *European Societies*, 21(5):753–773.
- Thomas Haider, Steffen Eger, Evgeny Kim, Roman Klinger, and Winfried Menninghaus. 2020. [PO-EMO: Conceptualization, annotation, and modeling of aesthetic emotions in German and English poetry](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1652–1663, Marseille, France. European Language Resources Association.
- D. Heerwegh. 2009. [Mode differences between face-to-face and web surveys: An experimental investigation of data quality and social desirability effects](#). *International Journal of Public Opinion Research*, 21(1):111–121.
- Christiane Heimann, Sandra Müller, Hannes Schammann, and Janina Stürner. 2019. [Challenging the nation-state from within: The emergence of trans-municipal solidarity in the course of the eu refugee controversy](#). *Social Inclusion*, 7(2):208–218.
- C. Hutto and Eric Gilbert. 2014. [Vader: A parsimonious rule-based model for sentiment analysis of social media text](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1).
- Alexander L. Janus. 2010. [The influence of social desirability pressures on expressed immigration attitudes*](#). *Social Science Quarterly*, 91(4):928–946.
- Sebastian Koos and Verena Seibel. 2019. [Solidarity with refugees across europe. a comparative analysis of public support for helping forced migrants](#). *European Societies*, 21(5):704–728.
- Anabel Kuntz, Eldad Davidov, and Moshe Semyonov. 2017. [The dynamic relations between economic conditions and anti-immigrant sentiment: A natural experiment in times of the european economic crisis](#). *International Journal of Comparative Sociology*, 58(5):392–415.
- Christian Lahusen and Maria T. Grasso, editors. 2018. [Solidarity in Europe: Citizens’ Responses in Times of Crisis](#). Palgrave Studies in European Political Sociology. Springer International Publishing.
- Drew Margolin and Wang Liao. 2018. [The emotional antecedents of solidarity in social media crowds](#). *New Media & Society*, 20(10):3700–3719.
- Adina Nerghes and Ju-Sung Lee. 2019. [Narratives of the refugee crisis: A comparative study of mainstream-media and twitter](#). *Media and Communication*, 7(2 Refugee Crises Disclosed):275–288.
- Francesco Nicoli. 2017. [Hard-line euroscepticism and the eurocrisis: Evidence from a panel study of 108 elections across europe](#). *JCMS: Journal of Common Market Studies*, 55(2):312–331.
- Laura Ana Maria Oberländer and Roman Klinger. 2018. [An analysis of annotated corpora for emotion classification in text](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2104–2119.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. [“Why should I trust you?”: Explaining the predictions of any classifier](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. Association for Computing Machinery.
- Sashank Santhanam, Vidhushini Srinivasan, Shaina Glass, and Samira Shaikh. 2019. [I stand with you: Using emojis to study solidarity in crisis events](#). *arXiv preprint*.
- Hilary Silver. 1994. [Social exclusion and social solidarity: three paradigms](#). *International Labour Review*, 133(5-6):531–578.
- Cass R. Sunstein. 2018. [#Republic: Divided democracy in the age of social media](#). Princeton University Press.
- Zeynep Tufekci. 2014. [Social movements and governments in the digital age: Evaluating a complex landscape](#). *Journal of International Affairs*, 68(1):1–18.
- Stefan Wallaschek. 2017. [Notions of solidarity in europe’s migration crisis: The case of germany’s media discourse](#).
- Stefan Wallaschek. 2019. [Solidarity in europe in times of crisis](#). *Journal of European Integration*, 41(2):257–263.
- Stefan Wallaschek. 2020a. [Contested solidarity in the euro crisis and europe’s migration crisis: A discourse network analysis](#). *Journal of European Public Policy*, 27(7):1034–1053.
- Stefan Wallaschek. 2020b. [The discursive construction of solidarity: Analysing public claims in europe’s migration crisis](#). *Political Studies*, 68(1):74–92.
- Wei Zhao, Steffen Eger, Johannes Bjerva, and Isabelle Augenstein. 2020a. [Inducing language-agnostic multilingual representations](#). *arXiv preprint arXiv:2008.09112*.
- Wei Zhao, Goran Glavaš, Maxime Peyrard, Yang Gao, Robert West, and Steffen Eger. 2020b. [On the limitations of cross-lingual encoders as exposed by reference-free machine translation evaluation](#). In

Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 1656–1671, Online. Association for Computational Linguistics.

A Appendices

Guidelines

Read the guidelines for annotating solidarity carefully.

- Definition of solidarity:

The preparedness to share one's own resources with others, be that directly by donating money or time in support of others or indirectly by supporting the state to reallocate and redistribute some of the funds gathered through taxes or contributions (Lahusen and Grasso, 2018)

- General rules

1. Do not take links (urls) into account when annotating.
2. Hashtags should be taken into account, especially if a tweet is otherwise neutral.
3. Emojis, if easily interpretable, can be taken into account.
4. If solidarity and anti-solidarity hashtags are used, code anti-solidarity.
5. When annotating use the scheme: Solidarity: 0, Anti-Solidarity: 1, Ambivalent: 2, Not Applicable: 3.

- Detailed rules for annotation.

1. A tweet is annotated as showing **solidarity**, when:
 - (a) It clearly indicates support people and the willingness to share resources and/or help.
 - (b) Positive attitude and gratitude to those sharing resources and/or helping.
 - (c) Advocacy of the European Union as a solidarity union.
 - (d) Criticism of the EU in terms of not doing enough to share resources and/or help others.
 - (e) Hashtags can be taken into account as to whether a tweet qualifies as showing solidarity (e.g. using hashtags like *#refugeeswelcome*).
 - (f) Hashtags should be taken into account if the tweet points neither

towards solidarity or anti-solidarity itself.

2. A tweet is annotated as showing **anti-solidarity**, when:

- (a) It clearly indicates no willingness to support people and an unwillingness to share resources and/or help.
- (b) It suggests to exclude our target groups from resources they currently have access to.
- (c) Tendencies of nationalism and closing borders.
- (d) Irony/sarcasm in tweets need to be taken into account.
- (e) Hashtags can be taken into account as to whether a tweet qualifies as anti-solidarity (e.g. using hashtags like *#g Brexit*).
- (f) Hashtags should be taken into account if the tweet points neither towards solidarity or anti-solidarity itself.

3. A tweet is annotated **ambivalent**, when:

- (a) The tweet shows solidarity or anti-solidarity sentiment, but it cannot be determined whether the tweet shows solidarity or anti-solidarity as there is additional info missing.
- (b) Even if taking hashtags into account, there is no clear indication as to whether the author shows solidarity or anti-solidarity.
- (c) If solidarity and anti-solidarity hashtags are used, code anti-solidarity.

4. A tweet is annotated **not applicable**, when:

- (a) There is no indication of solidarity or anti-solidarity sentiment in the tweet.
- (b) Even when hashtags that usually point towards solidarity or anti-solidarity are taken into account the tweet does not indicate any connection to solidarity or anti-solidarity.
- (c) The tweet concerns completely different topics than solidarity or anti-solidarity.
- (d) The tweet is not understandable (e.g.

contains only links).

Hashtags

Refugee crisis hashtags	Finance crisis hashtags
#asylumseeker	#austerity + #eu
#asylumseekers	#austerity + #euro
#asylkrise	#austerity + #eurobonds
#asylrecht	#austerity + #europe
#asylverfahren	#austerity + #eurozone
#flüchtling	#austerität + #eu
#flüchtlinge	#austerität + #euro
#flüchtlingsskise	#austerität + #eurobonds
#flüchtlingsswelle	#austerität + #europa
#leavenoonebehind	#austerität + #eurozone
#migrationskrise	#debtunion
#niewieder2015	#eurobonds
#opentheborders	#eurocrisis
#refugee	#eurokrise
#refugees	#eusolidarity
#refugeecrisis	#eusolidarität
#refugeesnotwelcome	#exiteu
#refugeeswelcome	#fiscalunion
#rightofasylum	#fiskalunion
#remigration	#schuldenunion
#seenotrettung	#transferunion
#standwithrefugees	
#wirhabenplatz	
#wirschaffendas	

Table 6: Hashtags used in our experiments.

Infection numbers vs. anti-solidarity Tweets

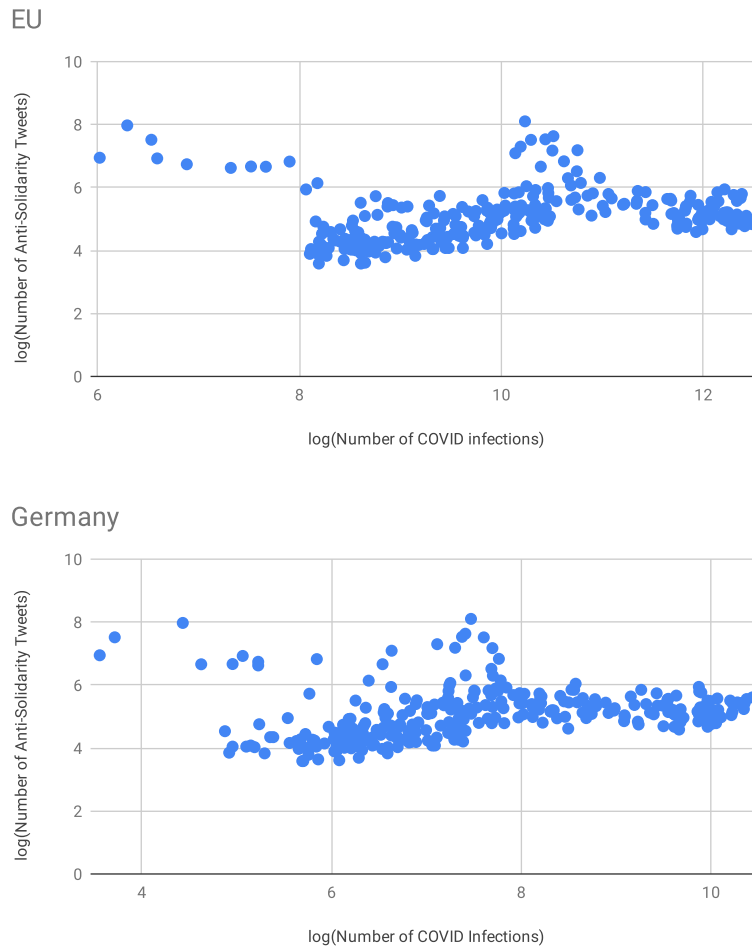


Figure 4: Scatter plot between infection rates and number of anti-solidarity tweets. Top: EU. Bottom: Germany. The time frame under consideration is 01.03.2020 to 14.12.2020 based on the data from <https://www.ecdc.europa.eu/en/publications-data/download-todays-data-geographic-distribution-covid-19-cases-worldwide>, restricted to 28 EU countries.