

# Investigating label suggestions for opinion mining in German Covid-19 social media

Tilman Beck<sup>1</sup>, Ji-Ung Lee<sup>1</sup>, Christina Viehmann<sup>2</sup>,  
Marcus Maurer<sup>2</sup>, Oliver Quiring<sup>2</sup> and Iryna Gurevych<sup>1</sup>

<sup>1</sup>Ubiquitous Knowledge Processing Lab, Technical University of Darmstadt

<sup>2</sup>Institut für Publizistik, Johannes Gutenberg-University Mainz

[www.ukp.tu-darmstadt.de](http://www.ukp.tu-darmstadt.de)

[www.ifp.uni-mainz.de](http://www.ifp.uni-mainz.de)

## Abstract

This work investigates the use of interactively updated label suggestions to improve upon the efficiency of gathering annotations on the task of opinion mining in German Covid-19 social media data. We develop guidelines to conduct a controlled annotation study with social science students and find that suggestions from a model trained on a small, expert-annotated dataset already lead to a substantial improvement – in terms of inter-annotator agreement (+.14 Fleiss'  $\kappa$ ) and annotation quality – compared to students that do not receive any label suggestions. We further find that label suggestions from interactively trained models do not lead to an improvement over suggestions from a static model. Nonetheless, our analysis of suggestion bias shows that annotators remain capable of reflecting upon the suggested label in general. Finally, we confirm the quality of the annotated data in transfer learning experiments between different annotator groups. To facilitate further research in opinion mining on social media data, we release our collected data consisting of 200 expert and 2,785 student annotations.<sup>1</sup>

## 1 Introduction

The impact analysis of major events like the Covid-19 pandemic is fundamental to research in social sciences. To enable more socially sensitive public decision making, researchers need to reliably monitor how various social groups (e.g., political actors, news media, citizens) communicate about political decisions (Jungherr, 2015). The increasing use of social media especially allows social science researchers to conduct opinion analysis on a larger scale than with traditional methods, e.g.

interviews or questionnaires. However, the publication of research results is often delayed or temporally transient due to limitations of traditional social science research, i.e. prolonged data gathering processes or opinion surveys being subject to reactivity. Given the increasing performance of language models trained on large amounts of data in a self-supervised manner (Devlin et al., 2019; Brown et al., 2020), one fundamental question that arises is how NLP systems can contribute to alleviate existing difficulties in studies for digital humanities and social sciences (Risch et al., 2019).

One important approach to make data annotation more efficient is the use of automated *label suggestions*. In contrast to *active learning*, that aims to identify a subset of annotated data which leads to optimal model training, label suggestions alleviate the annotation process by providing annotators with pre-annotations (i.e., predictions) from a model (Ringger et al., 2008; Schulz et al., 2019). To enable the annotation of large amounts of data which are used for quantitative analysis by disciplines such as social sciences, label suggestions are a more viable solution than active learning.

One major difficulty with label suggestions is the danger of biasing annotators towards (possibly erroneous) suggestions. So far, researchers have investigated automated label suggestions for tasks that require domain-specific knowledge (Fort and Sagot, 2010; Yimam et al., 2013; Schulz et al., 2019); and have shown that domain experts successfully identify erroneous suggestions and are more robust to potential biases. However, the limited availability of such expert annotators restricts the use of label suggestions to small, focused annotation studies. For tasks that do not require domain-specific knowledge and can be conducted with non-expert annotators – such as crowd workers or citizen science volunteers – on a large scale, label suggestions have not been considered yet. This leads to

<sup>1</sup>Code and data can be found on GitHub:  
<https://github.com/UKPLab/ac12021-label-suggestions-german-covid19>

two open questions. First, if non-expert annotators that do not receive any training besides annotation guidelines benefit from label suggestions at all. Second, if existing biases are amplified especially when including interactively updated suggestions that have been shown to be advantageous over static ones (Klie et al., 2020).

We tackle these challenges by conducting a comparative annotation study with social science students using a recent state-of-the-art model to generate label suggestions (Devlin et al., 2019). Our results show that a small set of expert-labeled data is sufficient to improve annotation quality for non-expert annotators. In contrast to Schulz et al. (2019), we show that although interactive and non-interactive label suggestions substantially improve the agreement, we do not observe significant differences between both approaches. We further confirm this observation with experiments using models trained on (and transferred to) individual annotator groups. Our contributions are:

- C1:** An evaluation of label suggestions in terms of annotation quality for non-expert annotators.
- C2:** An investigation of label suggestion bias for both static and interactively updated suggestions.
- C3:** A novel corpus of German Twitter posts that can be used by social science researchers to study the effects of governmental measures against Covid-19 on the public opinion.

Finally, we also publish 200 expert and 2,785 individual student annotations of our dataset to facilitate further research in this direction.

## 2 Related Work

**Label suggestions.** In an early work, Rehbein et al. (2009) study the effects of label suggestions on the task of word sense disambiguation and observe a positive effect on annotation quality. With the introduction of annotation tools such as brat (Stenetorp et al., 2012), WebAnno (Yimam et al., 2013), or INCEpTION (Klie et al., 2018), the use of label suggestions became more feasible; leading to an increased investigation of label suggestions in the context of NLP. For instance, Yimam et al. (2014) investigate label suggestions for Amharic POS tagging and German named entity recognition and show with expert annotators

that label suggestions significantly reduce the annotation time. Other works further investigate interactively updated label suggestions and come to a similar conclusion (Klie et al., 2020). Label suggestions have also been shown to be effective in non-NLP annotation tasks that require domain-specific knowledge such as in medical (Lingren et al., 2014) or educational (Schulz et al., 2019) use cases.

**Bias.** Annotations from untrained human annotators may introduce biases that are conveyed to machine learning models (Gururangan et al., 2018). One possible source of bias may be due to the different decision making process triggered by label suggestions – namely, first deciding if the suggested label is correct and only if not, considering different labels (Turner and Schley, 2016). Hence, the key question that arises is to what extent annotators are influenced by such suggestions. Although Fort and Sagot (2010) identify an influence on annotation behaviour when providing pre-annotated data for POS-tagging, they do not measure any clear bias in the annotated labels. Rosset et al. (2013) come to a similar conclusion when investigating the bias introduced by label suggestions in a cross-domain setup, i.e., when using label suggestions from a model that is trained on data from a different domain than the annotated data. They conduct their experiments with eight annotators from varying levels of expertise and report considerable annotation performance gains while not finding considerable biases introduced by label suggestions. Most similar to our work is the setup from Schulz et al. (2019). The authors investigate interactive label suggestions for expert annotators across two domains and study the effects of using existing and newly annotated data for training different suggestion models. They compare personalised user models against a universal model which has access to all annotated data and show that the latter provides suggestions with a higher acceptance rate. This seems less surprising due to the substantially larger training set. Further, they do not identify any bias introduced by pre-annotating data.

Whereas existing work reports no measurable bias for expert annotators (Fort and Sagot, 2010; Lingren et al., 2014; Schulz et al., 2019), it remains unclear for annotators who have no prior experience in similar annotation tasks; especially for scenarios where – besides annotation guidelines – no further training is provided. However, the use of novice annotators is common for sce-

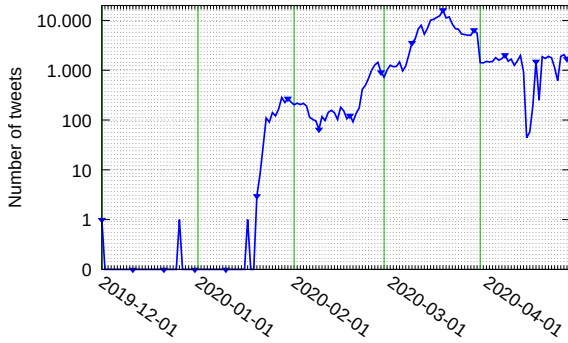


Figure 1: Number of tweets per day collected from December 2019 to April 2020.

narios where no linguistic or domain expertise is required. Hence, we present a first case-study for the use of interactive label suggestions with non-expert annotators. Furthermore, we find that recent state-of-the-art models such as BERT (Devlin et al., 2019) can provide high-quality label suggestions with already little training data and hence, are important for interactive label suggestions in non-expert annotation tasks.

### 3 Annotation Task

Our task is inspired by social science research on analyzing public opinion using social media (Jungherr, 2015; McCormick et al., 2017). The goal is to identify opinions in German-speaking countries about governmental measures established to contain the spread of the Corona virus. We use Twitter due to its international and widespread usage that ensures a sufficient database and the several challenges for the automatic identification of opinions and stance it poses from an NLP perspective (Imran et al., 2016; Mohammad et al., 2016; Gorrell et al., 2019; Conforti et al., 2020). For example, the use of language varies from colloquial expressions to well-formed arguments and news-spreading statements due to its heterogeneous user base. Additionally, hashtags are used directly as part of text but also to embed the tweet itself in the broader discussion on the platform. Finally, the classification of a tweet is particularly challenging given the character limitation of the platform, i.e., at the date of writing Twitter allows for 280 characters per tweet.

**Data collection.** Initially, we collected tweets from December 2019 to the end of April 2020. Using a manually chosen set of search queries (‘corona’, ‘pandemie’, ‘covid’, ‘socialdistance’),

we made use of the Twitter Streaming API and gathered only those tweets which were classified as German by the Twitter language identifier. This resulted in a set of approximately 16.5 million tweets. We retained only tweets that contain key terms referring to measures related to the Covid-19 pandemic and removed all duplicates, retweets and all tweets with text length less than 30 characters. After filtering, 237,616 tweets remained and their daily temporal distribution is visualized in Figure 1. We sample uniformly at random from the remaining tweets for all subsequent annotation tasks.<sup>2</sup>

**Annotation scheme.** We developed annotation guidelines together with three German-speaking researchers from social sciences and iteratively refined them in three successive rounds. Our goal from a social science perspective is to analyze the public perception of measures taken by the government. Therefore, the resulting dataset should help in (1) identifying relevant tweets for governmental measures and if relevant, (2) detecting what stance is expressed. We follow recent works on stance detection and Twitter data (Hanselowski et al., 2018; Baly et al., 2018; Conforti et al., 2020) and use four distinct categories for our annotation. They are defined as follows:

- Unrelated:** no measures related to the containment of the pandemic are mentioned
- Comment:** measures are mentioned, but not assessed or neutral
- Support:** measures are assessed positively
- Refute:** measures are assessed negatively

The four label annotation scheme allows us to distinguish texts that are related to the pandemic but do not talk about measures (i.e., unrelated).

### 4 Study Setup

Our goal is to study the effects of interactively updated and static label suggestions in non-expert annotation scenarios. Non-experts such as crowd workers or student volunteers have no prior experience in annotating comparable tasks and only receive annotation guidelines for preparation.<sup>3</sup> Our secondary goal is to collect a novel dataset that can be used by social science researchers to study the

<sup>2</sup>We provide additional information about data collection in Appendix A and discuss ethical concerns regarding the use of Twitter data after the conclusion.

<sup>3</sup>We provide the original German guidelines along with the dataset. An English summary is provided in the Appendix B

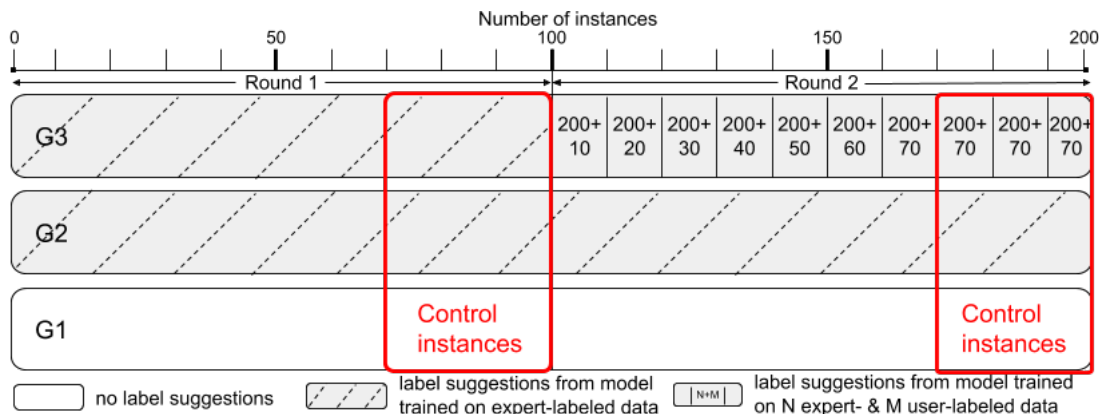


Figure 2: Design of the annotation setup for each of the three user groups. The 30 quality control instances (red) were inserted at random positions but are visualized at the end for presentation purpose.

effects of governmental measures for preventing the spread of Covid-19 on the public opinion.

To train a model that provides label suggestions to our non-expert annotators, we first collect a small set of 200 expert-annotated instances. We then split our non-expert annotators into three different groups that receive (G1) no label suggestions, (G2) suggestions from a model trained on expert annotations, and (G3) suggestions from a model that is retrained interactively using both expert-annotated and interactively annotated data.

#### 4.1 Expert Annotations

The expert annotations were provided by the researchers (three social science researchers and one NLP researcher) that created the annotation guidelines and who are proficient in solving the task. In total, 200 tweets were sampled uniformly at random and annotated by all four experts. The inter-annotator agreement (IAA) across all 200 tweets lies at 0.54 Fleiss’s  $\kappa$  (moderate agreement) and is comparable to previously reported annotation scores in the field of opinion and argument mining (Bar-Haim et al., 2020; Schaefer and Stede, 2020; Boltužić and Šnajder, 2014). Overall, in more than 50% of the tweets all four experts selected the same label (respectively, in  $\sim 75\%$  of the tweets at least three experts selected the same label). The disagreement on the remaining  $\sim 25\%$  of the tweets furthermore shows the increased difficulty of our task due to ambiguities in the data source, e.g., ironical statements or differentiating *governmental* measures from *non-governmental* ones like home-office. To compile gold standard labels for instances that the experts disagreed upon, we apply MACE (Hovy et al., 2013) using a threshold of 1.0.

The resulting labels were then re-evaluated by the experts and agreed upon.

#### 4.2 Student Annotations

The annotations were conducted with a group of 21 German-speaking university students. To ensure a basic level of comparability for our student annotators, we recruited all volunteers from the same social science course at the same university. The annotators received no further training apart from the annotation guidelines. We randomly assigned them to three different groups (G1, G2, and G3), each consisting of seven students. To investigate the effects of interactive label suggestions, we defined different annotation setups for each group. The annotations were split into two rounds. At each round of annotation, students were provided with 100 tweets consisting of 70 new tweets and 30 quality control tweets from the expert-labeled data which are used to compare individual groups. Across both rounds, we thus obtain a total of 140 unique annotated tweets per student and use 60 tweets for evaluation. The annotation setup of each group including the individual data splits is visualized in Figure 2.<sup>4</sup>

**No label suggestions (G1).** The first group serves as a control group and receives no label suggestions.

**Static label suggestions (G2).** The second group only receives label suggestions based on a model which was trained using the 200 expert-labeled instances described in section 4.1.

<sup>4</sup>Note that the control instances were distributed uniformly at random within a round to mitigate any interdependency effects between different tweets.



**Interactive label suggestions (G3).** The last group of students receives expert label suggestions in the first round and interactively updated label suggestions in the second round. In contrast to existing work (Schulz et al., 2019), this setup allows us to directly quantify effects of bias amplification that may occur with interactive label suggestions.

### 4.3 Label Suggestion Model

**System setup.** We conduct our annotation experiments using INCEPTION (Klie et al., 2018) which allows us to integrate label suggestions using recommendation models. To obtain label suggestions, we use a German version of BERT (Ger-BERT) that is available through the HuggingFace library (Wolf et al., 2020).<sup>5</sup> We perform a random hyperparameter search (cf. Appendix B.3) and train the model on the expert annotated data for 10 epochs with a learning rate of 8e-5 and a batch size of 8. We select the model that performed best in terms of F1-score on a held-out stratified test set (20% of the data) across ten runs with different random seeds. All experiments were conducted on a desktop machine with a 6-core 3.8 GHz CPU and a GeForce RTX 2060 GPU (8GB).

Model	Macro-F1	Accuracy
Majority	.15	.45
Random	.23	.27
BiLSTM (Schulz et al., 2019)	.47	.53
SBERT+LGBM (Klie et al., 2020)	.50	.55
Ger-BERT (this work)	.66	.68

Table 1: Performance of various label suggestion models on expert-labeled dataset.

**Model comparison.** To assess the label suggestion quality of our model, we report the predictive performance on the expert-labeled dataset (setup as described above) in Table 1. We compare our model with baselines<sup>6</sup> which have been used in related work (Schulz et al., 2019; Klie et al., 2020) for label suggestions. As expected, Ger-BERT achieves superior performance and the results are promising for using label suggestions.

**Interactive training routine.** To remedy the cold-start problem, G3 receives label suggestions from the model trained only on the expert-annotated data in round 1. Afterwards, we retrain the model with an increasing number of instances

using both, the expert annotations and the G3 data of individual students from round 1.<sup>7</sup> To avoid unnecessary waiting times for our annotators due to the additional training routine, we always collect batches of 10 instances before re-training our model. We then repeatedly train individual models for each student in G3 with an increasing amount of data of up to 70 instances. The 30 expert-annotated quality control tweets were excluded in this step to avoid conflicting labels and duplicated data.

## 5 Study Evaluation

Table 2 shows the overall statistics of our resulting corpus consisting of 200 expert and 2,785 student-annotated German tweets. Note that we removed 60 expert-annotated instances that we included for annotation quality control for each student, resulting in 140 annotated tweets per student.

**Outliers.** A fine-grained analysis of annotation time is not possible due to online annotations at home. However, one student in G3 had, on average, spent less than a second for each annotation and accepted almost all suggested labels. This student’s annotations were removed from the final dataset and assumed as faulty labels considering the short amount of time spent on this task in comparison to the minimum amount of seven seconds per tweet and annotation for all other students.

### 5.1 Annotation Quality

To assess the overall quality of our collected student annotations, we investigate annotator consistency in terms of inter-annotator-agreement (IAA) as well as the annotator accuracy on our quality assurance instances.

Table 3 shows Fleiss’  $\kappa$  (Fleiss, 1971) and the accuracy computed for the quality control instances that were consistent across all groups. In general, we observe a similar or higher agreement for our students compared to the expert annotations ( $\kappa = 0.54$ ) showing that the guidelines were able to convey the task well. We also find that groups that receive label suggestions (G2 and G3) achieve a substantially larger IAA as opposed to G1. Most interestingly, we observe a substantial increase in IAA for both G2 and G3 in the second annotation round, whereas the IAA in G1 remains stable.

<sup>7</sup>Note that using all previously annotated data of G3 would impair the comparability between individual students as the data was collected asynchronously to allow students to pick their best suited timeslot. Further, a synchronization step between users would impair the applicability of the approach.

<sup>5</sup><https://deepset.ai/german-bert>

<sup>6</sup>We adapted the respective architectures to our setup.

N	Annotator	Avg. Length	Unrelated	Comment	Support	Refute
200	Expert	189 ( $\pm 75$ )	53 (26.5%)	89 (44.5%)	43 (21.5%)	15 (7.5%)
2,785	Student	185 ( $\pm 75$ )	1,003 (36.0%)	1,055 (37.9%)	425 (15.3%)	302 (10.8%)
965	G1	185 ( $\pm 76$ )	387 (40.1%)	334 (34.6%)	128 (13.3%)	116 (12.0%)
980	G2	185 ( $\pm 73$ )	320 (32.7%)	407 (41.5%)	152 (15.5%)	101 (10.3%)
840	G3	184 ( $\pm 75$ )	296 (35.2%)	314 (37.4%)	145 (17.3%)	85 (10.1%)

Table 2: Our Twitter dataset on public opinion about containment measures during the Corona pandemic.

	G1		G2		G3	
	Acc	IAA	Acc	IAA	Acc	IAA
Round 1	.74	.48	.90	.76	.84	.62
Round 2	.68	.47	.92	.81	.82	.67
Total	.71	.48	.91	.78	.83	.65

Table 3: Annotation accuracy (Acc) and IAA (Fleiss’  $\kappa$ ) on the quality control instances for each annotator group and round.

Analyzing our models’ predictions shows that the suggested labels for the 60 quality control samples mostly conform with the label given by the expert (97% for G2 and 94% for G3). Therefore, annotators are inclined to accept the label suggested by the model. We can further confirm this observation when investigating the number of instances that the students labeled correctly (accuracy). The highest accuracy is observed for the group that received the highest quality suggestions (G2). Furthermore, both groups that received label suggestions (G2, G3) express an increased accuracy over the control group (G1). In general, for both rounds the accuracy remains similarly high across all groups ( $\pm .02$  difference) with only a slight decrease ( $-.04$ ) for G1. Hence, we conjecture that the resulting annotations provide satisfying quality given the challenging task and annotator proficiency.

## 5.2 Suggestion Bias

One major challenge in using label suggestions is known in psychology as the *anchoring effect* (Tversky and Kahneman, 1974; Turner and Schley, 2016). It describes the concept that annotators who are provided a label suggestion follow a different decision process compared to a group that does not receive any suggestions and tend to accept the suggestions. As we observe larger IAA and accuracy for groups receiving label suggestions, we look at the label suggestion acceptance rate and which

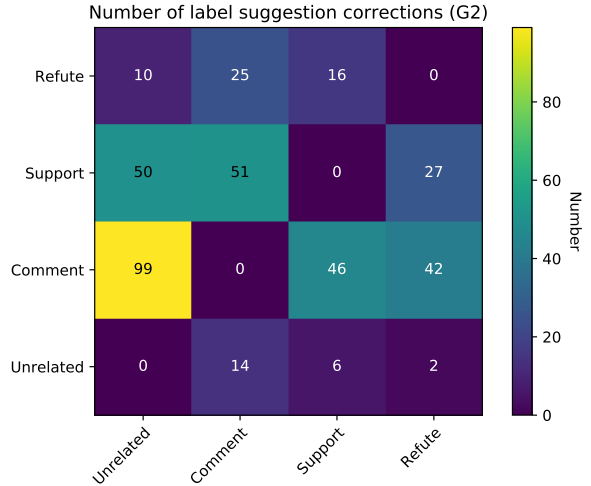


Figure 3: The number of rejected label suggestions. The x-axis displays the corrected label and the y-axis the label suggestion. For example, the upper left corner shows that ten suggestions of label Refute were corrected as Unrelated by the users.

labels have been corrected by the annotators.

**Acceptance rate.** One way to quantify possible biases is to evaluate if annotators tend to accept more suggestions with an increasing number of instances (Schulz et al., 2019). This may be the case when annotators increasingly trust the model with consistently good suggestions. Consequently, with increasing trust towards the model’s predictions, non-expert annotators may tend to accept more model errors. To investigate if annotators remain capable of reflecting on instance and label suggestion, we compute the average acceptance rate for G2 and G3 in both rounds. We find that for both groups, the acceptance rate remains stable (G2: 73% and 72%, G3: 68% and 69%) and conclude that annotators receiving high quality label suggestions remain critical while producing more consistent results.

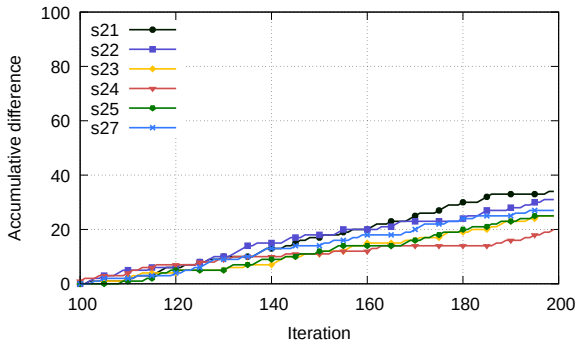


Figure 4: Number of label suggestions that diverge from the model trained on expert-data with increasing number of annotations (for each student).

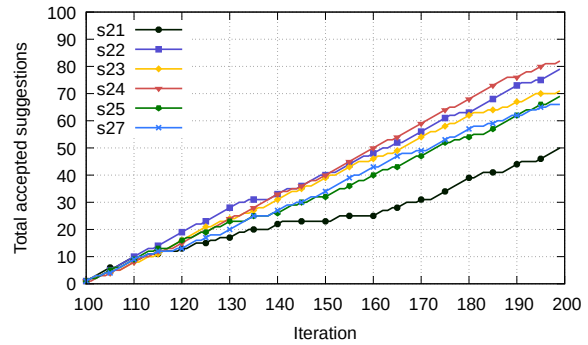


Figure 5: Number of label suggestions in G3 that have been accepted in the second round of annotations (for each student).

**Label corrections.** To further evaluate if students are vulnerable to erroneous label suggestions from a model, we specifically investigate labels that have been corrected. Figure 3 shows our results for G2.<sup>8</sup> As can be seen, the most notable number of label corrections were made by students for unrelated tweets that were classified as comments by the model. Additionally, we find a large number of corrections that have been made with respect to the stance of the presented tweet. We will discuss both types of corrections in the following.

*Unrelated tweets.* The label suggestion model makes the most errors for unrelated tweets (i.e., tweets that are corrected as *Unrelated*) by misclassifying them as *Comment* (99). In contrast, instances that are identified as *Unrelated* tweets are only seldomly corrected. This indicates an increased focus on recall at the expense of precision for related tweets, most likely due to *Comment* being the largest class in the training data (see Table 2, expert data). We find possible causes for such wrong predictions when we look at examples where *Comment* was suggested for *Unrelated* instances<sup>9</sup>:

**Example 1:** The corona virus also requires special protective measures for Naomi Campbell. The top model wears a protective suit during a trip.

**Example 2:** Extraordinary times call for extraordinary measures: the "Elbschlosskeller" now has a functioning door lock. #Hamburg #Corona #COVID-19

<sup>8</sup>Note that analyzing G3 shows similar observations (cf. Appendix C).

<sup>9</sup>Note that we present translations of the original German texts for better readability and to protect user privacy

Clearly, these examples are fairly easy to annotate for humans but are difficult to predict for a model due to specific cue words being mentioned, e.g., *measures*. Similar results have also been reported in previous work (Hanselowski et al., 2018; Conforti et al., 2020).

*Stance.* In Figure 3, we can also see that the model makes mistakes regarding the stance of a tweet. Especially, 101 *Support* suggestions have been corrected as either being *unrelated* or *neutral* and 88 *Comment* suggestions have been corrected to either *Support* or *Refute*. For the second case, we often discover tweets that implicitly indicate the stance – for example, by complaining about people ignoring the measures:

**Example 3:** Small tweet aside from XR: Colleague drags himself into the office this morning with flu symptoms (ÖD) The other colleagues still have to convince him to please go home immediately. Only then does he show some understanding. Unbelievable. #COVID #SocialDistancing

Such examples demonstrate the difficulty of the task and seem to be difficult to recognize for the model. However, given the large amount of label corrections, the non-expert annotators seem to be less susceptible to accept such model errors.

### 5.3 Bias Amplification

The high number of label corrections for specific types of tweets shows that our annotators of G2 remained critical towards the suggested label. With interactively updated suggestions however, this may not be the case. Especially annotators that accept erroneous suggestions may lead to reinforc-

ing a model in its prediction; hence, leading to amplifying biases.

**Diverging suggestions.** To study such effects, we first identify if the interactively updated models express a difference in terms of predictions compared to the static model. In Figure 4 we can observe that with already 40 instances (Iteration 140), the number of differently predicted instances is ten or higher across all personalized models. This divergence is highly correlated with the number of changes a student provides (see Figure 5). We thus can conclude that the interactively trained models are able to adapt to the individual annotations for each annotator.

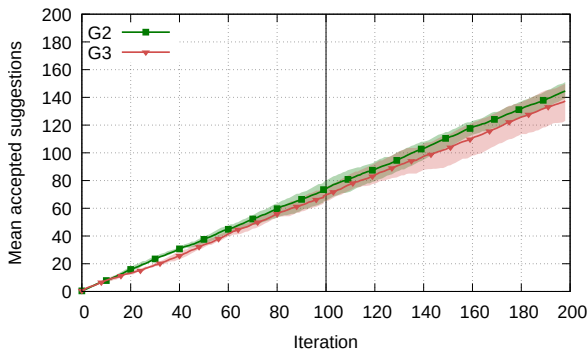


Figure 6: Average number of accepted label suggestions across all instances for G2 and G3. The shared areas display the upper and lower quartile for each group.

**Comparison to G2.** Figure 6 shows the average number of accepted suggestions for G2 and G3 as well as the upper and lower quartiles, respectively. The vertical line separates the first and the second round of annotations. We find that especially in the first round of annotations, both groups have a very similar acceptance rate of suggested labels. Only with interactively updated suggestions we find an increasing divergence in G3 with respect to the upper and lower quartiles.

**Individual acceptance rate.** To assess the impact of interactive label suggestions, we further investigate how many suggestions were accepted by each annotator. Figure 5 shows the number of accepted label suggestions for each student in G3 in the second round of annotations. Although we observe that the average number of accepted label suggestions remains constant across G2 and G3, we can see substantial differences between individual students. For instance, we can observe that for  $s_{21}$ , the increased model adaptivity leads

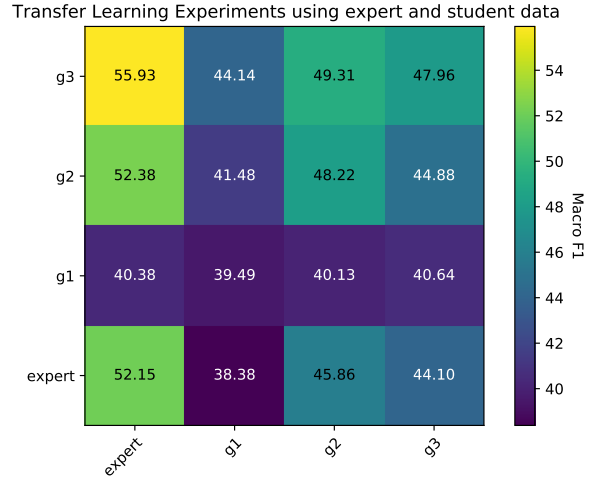


Figure 7: Transfer learning performance of models trained on individual annotator groups. The x-axis presents the dataset which is used for model training, the y-axis lists the dataset used for model testing.

to an overall decrease in the number of accepted labels. Moreover,  $s_{24}$  who received predictions that diverge less from the static model prediction accepted the most suggestions in the second round. This shows that interactive label suggestions does not necessarily lead to a larger acceptance rate – possibly amplifying biases – but instead, varies for each annotator and needs to be investigated in future work.

#### 5.4 Cross-group Transfer

Finally, we investigate how well models trained on different annotator groups transfer to each other. We hence conduct transfer learning experiments for which we remove the quality control instances in our student groups and train a separate Ger-BERT model using the same hyperparameters as for the expert model. We use 80% of the data for training and the remaining 20% to identify the best model which we then transfer to another group. Figure 7 shows the macro-F1 scores averaged across ten independent runs, diagonal entries are the scores on the 20%. Most notably, models trained on the groups with label suggestions (G2, G3) do in fact perform comparable or better on the expert-labeled data and outperform models trained on the group not receiving any suggestions (G1). The higher cross-group performance for models trained on groups that received label suggestions shows that the label suggestions successfully conveyed knowledge from the expert annotated data to our students.



## 6 Conclusion

In this work, we analysed the usefulness of providing label suggestions for untrained annotators to identify opinions in a challenging text domain (i.e., Twitter). We generated suggestions using expert-labeled training data as well as interactively training models using data annotated by untrained students. Our results show that label suggestions from a state-of-the-art sentence classification model trained on a small set of expert annotations help improving annotation quality for untrained annotators. In terms of potential biases that may occur with untrained annotators we observe that the students retained their capability to reflect on the suggested label. We furthermore do not observe a general amplification in terms of bias with interactively updated suggestions; however, we find that such effects are very specific to individual annotators. We hence conclude that interactively updated label suggestions need to be considered carefully when applied to non-expert annotation scenarios.

For future work, we plan to leverage our setup to annotate tweets from a larger time span. In Germany, the measures taken by the government have been met with divided public reaction – starting with reactions of solidarity and changing towards a more critical public opinion (Viehmann et al., 2020a,b). In particular, we are interested if our label suggestion model is robust enough to account for such a shift in label distribution.

### Acknowledgements

This work has been supported by the German Research Foundation (DFG) as part of the Research Training Group KRITIS No. GRK 2222/2. Further, this work is supported by the European Regional Development Fund (ERDF) and the Hessian State Chancellery – Hessian Minister of Digital Strategy and Development under the promotional reference 20005482 (TexPrax). We thank the student volunteers from the University of Mainz for their annotations as well as Johannes Daxenberger, Mohsen Mesgar, Ute Winchenbach, Kevin Stowe and the anonymous reviewers for their valuable feedback.

### Ethical considerations

**Data collection and annotation.** The tools we use to collect Tweets are in compliance with Twitter’s terms of service. We only release the set of identifiers (Tweet IDs) for the texts used in this

research project. Thereby, we adhere to the Twitter Developer policy<sup>10</sup> and give users full control of their privacy and data as they can delete or privatize tweets so that they cannot be collected.

We asked student annotators for voluntary participation in the annotation study. All students have been informed about the goal of the conducted research and the purpose of the collected annotations. During annotation no information about the tweet’s author or any other additional metadata was made available to the annotators. We did not collect any personal data from the students before, after, or during the annotation task.

**Data usage.** This work presents an investigation of efficient data annotation methods in a case study on social media data. The results of this work allow social science researchers to apply their analysis on a larger scale. In the case of analyzing public opinion on governmental measures, the resulting analysis allows politicians to make more socially sensitive public decisions. This information is useful in aggregated form, without the need for information about individual users. However, we want to point out that users of social media (particularly Twitter) do not constitute a representative sample of the general population, especially in Germany (Newman et al., 2020). Therefore, our goal is not to foster public decision-making solely based upon analysis of Twitter but to provide an additional supporting tool.

**Dual use.** Further, we acknowledge the potential of misuse of our dataset: the annotated data allows anyone, including both individuals and organizations, for training models to identify individuals expressing their consent or dissent with governmental actions. To this end, we follow the argumentation by (Benton et al., 2017) that in general we cannot prevent publicly available data from being misused but we want to make both researchers and the general public aware of the possible malicious use.

### References

Ramy Baly, Mitra Mohtarami, James Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018. *Integrating stance detection and fact checking in a unified corpus*. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human*

<sup>10</sup><https://developer.twitter.com/en/developer-terms/agreement-and-policy>

- Language Technologies, Volume 2 (Short Papers)*, pages 21–27, New Orleans, Louisiana. Association for Computational Linguistics.
- Roy Bar-Haim, Lilach Eden, Roni Friedman, Yoav Kantor, Dan Lahav, and Noam Slonim. 2020. [From arguments to key points: Towards automatic argument summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4029–4039, Online. Association for Computational Linguistics.
- Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017. [Ethical research protocols for social media health research](#). In *Proceedings of the First ACL Workshop on Ethics in Natural Language Processing*, pages 94–102, Valencia, Spain. Association for Computational Linguistics.
- Filip Boltužić and Jan Šnajder. 2014. [Back up your stance: Recognizing arguments in online discussions](#). In *Proceedings of the First Workshop on Argumentation Mining*, pages 49–58, Baltimore, Maryland. Association for Computational Linguistics.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *arXiv preprint arXiv:2005.14165*.
- Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd, and Nigel Collier. 2020. [Will-they-won't-they: A very large dataset for stance detection on Twitter](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1715–1724, 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.
- Joseph L Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological bulletin*, 76(5):378.
- Karèn Fort and Benoît Sagot. 2010. [Influence of pre-annotation on POS-tagged corpus development](#). In *Proceedings of the Fourth Linguistic Annotation Workshop*, pages 56–63, Uppsala, Sweden. Association for Computational Linguistics.
- Genevieve Gorrell, Elena Kochkina, Maria Liakata, Ahmet Aker, Arkaitz Zubiaga, Kalina Bontcheva, and Leon Derczynski. 2019. [SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours](#). In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 845–854, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Andreas Hanselowski, Avinesh PVS, Benjamin Schiller, Felix Caspelherr, Debanjan Chaudhuri, Christian M. Meyer, and Iryna Gurevych. 2018. [A retrospective analysis of the fake news challenge stance-detection task](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1859–1874, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Dirk Hovy, Taylor Berg-Kirkpatrick, Ashish Vaswani, and Eduard Hovy. 2013. [Learning whom to trust with MACE](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1120–1130, Atlanta, Georgia. Association for Computational Linguistics.
- Muhammad Imran, Prasenjit Mitra, and Carlos Castillo. 2016. [Twitter as a lifeline: Human-annotated Twitter corpora for NLP of crisis-related messages](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 1638–1643, Portorož, Slovenia. European Language Resources Association (ELRA).
- Andreas Jungherr. 2015. *Analyzing political communication with digital trace data*. Springer, Cham, Switzerland.
- Jan-Christoph Klie, Michael Bugert, Beto Boullosa, Richard Eckart de Castilho, and Iryna Gurevych. 2018. [The INCEpTION platform: Machine-assisted and knowledge-oriented interactive annotation](#). In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 5–9, Santa Fe, New Mexico. Association for Computational Linguistics.
- Jan-Christoph Klie, Richard Eckart de Castilho, and Iryna Gurevych. 2020. [From Zero to Hero: Human-In-The-Loop Entity Linking in Low Resource Domains](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6982–6993, Online. Association for Computational Linguistics.
- Todd Lingren, Louise Deleger, Katalin Molnar, Haijun Zhai, Jareen Meinzen-Derr, Megan Kaiser, Laura Stoutenborough, Qi Li, and Imre Solti. 2014. [Evaluating the impact of pre-annotation on annotation speed and potential bias: natural language processing gold standard development for clinical named](#)

- entity recognition in clinical trial announcements. *Journal of the American Medical Informatics Association*, 21(3):406–413.
- Tyler H McCormick, Hedwig Lee, Nina Cesare, Ali Shojaie, and Emma S Spiro. 2017. Using twitter for demographic and social science research: Tools for data collection and processing. *Sociological methods & research*, 46(3):390–421.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. A dataset for detecting stance in tweets. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3945–3952, Portorož, Slovenia. European Language Resources Association (ELRA).
- Nic Newman, Richard Fletcher, Anne Schulz, Simge Andi, and Rasmus-Kleis Nielsen. 2020. *Digital news report 2020*. Reuters Institute for the Study of Journalism, pages 2020–06.
- Ines Rehbein, Josef Ruppenhofer, and Caroline Sporleder. 2009. Assessing the benefits of partial automatic pre-labeling for frame-semantic annotation. In *Proceedings of the Third Linguistic Annotation Workshop (LAW III)*, pages 19–26, Suntec, Singapore. Association for Computational Linguistics.
- Eric Ringger, Marc Carmen, Robbie Haertel, Kevin Seppi, Deryle Lonsdale, Peter McClanahan, James Carroll, and Noel Ellison. 2008. Assessing the costs of machine-assisted corpus annotation through a user study. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).
- Julian Risch, Anke Stoll, Marc Ziegele, and Ralf Kretzel. 2019. hpidedis at germeval 2019: Offensive language identification using a german bert model. In *Proceedings of the 15th Conference on Natural Language Processing (KONVENS 2019)*, pages 405–410, Erlangen, Germany. German Society for Computational Linguistics & Language Technology.
- Sophie Rosset, Cyril Grouin, Thomas Lavergne, Mohamed Ben Jannet, Jérémy Leixa, Olivier Galibert, and Pierre Zweigenbaum. 2013. Automatic named entity pre-annotation for out-of-domain human annotation. In *Proceedings of the 7th Linguistic Annotation Workshop and Interoperability with Discourse*, pages 168–177, Sofia, Bulgaria. Association for Computational Linguistics.
- Robin Schaefer and Manfred Stede. 2020. Annotation and detection of arguments in tweets. In *Proceedings of the 7th Workshop on Argument Mining*, pages 53–58, Online. Association for Computational Linguistics.
- Claudia Schulz, Christian M. Meyer, Jan Kiesewetter, Michael Sailer, Elisabeth Bauer, Martin R. Fischer, Frank Fischer, and Iryna Gurevych. 2019. Analysis of automatic annotation suggestions for hard discourse-level tasks in expert domains. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2761–2772, Florence, Italy. Association for Computational Linguistics.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. brat: a web-based tool for NLP-assisted text annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107, Avignon, France. Association for Computational Linguistics.
- Brandon M Turner and Dan R Schley. 2016. The anchor integration model: A descriptive model of anchoring effects. *Cognitive Psychology*, 90:1–47.
- Amos Tversky and Daniel Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *science*, 185(4157):1124–1131.
- Christina Viehmann, Marc Ziegele, and Oliver Quiring. 2020a. Gut informiert durch die Pandemie? Nutzung unterschiedlicher Informationsquellen in der Corona-Krise. Ergebnisse einer dreiwöchigen Panelbefragung im Jahr 2020. *Media Perspektiven*, 10-11:2556–577.
- Christina Viehmann, Marc Ziegele, and Oliver Quiring. 2020b. In der Krise rücken alle etwas mehr zusammen? Fast! – Wie Medien und andere Faktoren zur gefühlten Integration und Spaltung beitragen. *Kommunikationsmanager*, 2:32–36.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Seid Muhie Yimam, Chris Biemann, Richard Eckart de Castilho, and Iryna Gurevych. 2014. Automatic annotation suggestions and custom annotation layers in WebAnno. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 91–96, Baltimore, Maryland. Association for Computational Linguistics.
- Seid Muhie Yimam, Iryna Gurevych, Richard Eckart de Castilho, and Chris Biemann. 2013. WebAnno: A flexible, web-based and visually supported system for distributed annotations. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 1–6, Sofia, Bulgaria. Association for Computational Linguistics.



## A Data Crawl Details

### A.1 Filter Terms

We crawled Twitter posts using Live Streaming API (Twitter). Based on a preliminary examination of the data, we selected the following set of terms for filtering tweets: ['stayhomesavelives', 'wirbleibenzuhause', 'bleibdaheim', 'abstandhalten', 'flatthecurve', 'flattenthecurve', 'sperre', 'verbot', 'beschraenkung', 'quarantäne', 'quarantaene', 'wirvsvirus', 'schließung', 'homeoffice', 'infektionsschutz', 'ansteckungsrisiko', 'notbetrieb', 'bleibtzuhause', 'stayhome']

### A.2 Additional Data

Unlike related work (Schaefer and Stede, 2020), we do not investigate reply structures of tweets. In preliminary experiments we found that our collection method provides a large enough amount of relevant tweets which can be annotated without context.

## B Annotation

In this section we provide more detailed information on the annotation guidelines and the annotation platform we used.

### B.1 Annotation Guidelines

We first provide background context about the measures for containing the spread of Covid-19. Afterwards, we provide a definition for measures to be considered during the annotation study: *We consider all measures which are taken by the government to contain the pandemic (e.g., closing of schools).*

Next, we introduce the annotation task as a two-step process. First, the annotator has to decide if the text actually does mention measures as defined as above (some examples are provided for clarification). If not, the annotator selects the label NoMeasure<sup>11</sup>. In the opposite case, the annotator decides in a second step if the text contains a positive position (ProOpinion), a negative position (ConOpinion) or if there is no stance expressed

<sup>11</sup>Please note that we use a different notation in the main paper. The label NoMeasure corresponds to Unrelated, label NoOpinion corresponds to Comment, ProOpinion corresponds to Support and ConOpinion corresponds to Refute

(NoOpinion). We provide examples for each label to our annotators.

During our preliminary studies, we identified several ambiguities regarding the stance annotation which is in the nature of the source of the texts (Twitter) and the subject of the annotation (measures regarding the Covid-19 pandemic):

- **a Tweet discusses (positive/negative) consequences or by-products of measures** : we regard those as (positive/negative) statements as the author implicitly states their opinion by reflecting upon the measures
- **a Tweet reflects the opinion of another actor**: this is considered as an opinion as defined above. It is assumed that the author posts this opinion because they identify themselves with the original opinion.
- **a Tweet makes an unagitated observation whether measures are functioning**: this is not to be taken as an opinion for or against the measures per se. Only if an explicit assessment of the observation is made, the position can be derived.
- **the role of Hashtags**: Hashtags are often ambiguous and the respective context needs to be taken into account. Therefore, in our annotation hashtags are only considered as context to what is said; they never stand for themselves. Hashtags can be used to determine whether a measure is being addressed. To do this, the hashtag must contain a measure. Further, hashtags can be used as context to support the position in a tweet.

These decisions are reflected at the corresponding positions in the annotation guidelines, along with several example tweets. In the end we provide a note that Twitter posts may contain malicious, suggestive, offensive, or potentially sensitive content and that the annotation can be paused and resumed at any time.

### B.2 Annotation Interface

In Figure 8 a screenshot of the annotation interface is depicted. It is taken from the group where label recommendations are provided. The Twitter posts to be annotated are shown in the center where each line corresponds to a single tweet. For the sake of clarity, only five texts are shown simultaneously



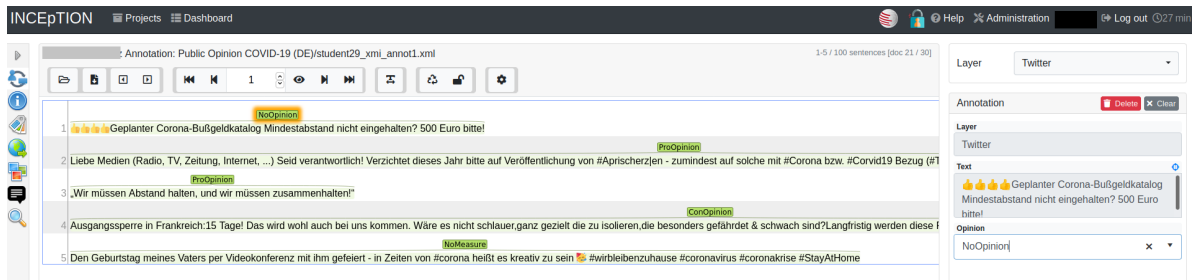


Figure 8: A screenshot of the annotation interface using INCEpTION (Klie et al., 2018)

and the user navigates through all texts using the navigation bar above the text window.

The label recommendations are displayed using a green box above the corresponding text and the currently selected recommendation is highlighted in orange. If the user agrees with the provided label, nothing needs to be changed. In the opposite case, the user can click on the recommendation and select another label on the right-hand side (Annotation panel) using the Opinion dropdown field. The annotators receiving no label suggestions (G1) do not see any recommendation during annotation. They create an annotation for each sentence by double-clicking on the sentence. Once the user has finished annotating all samples, the annotation session is finished by clicking the lock symbol in the navigation bar. The technical procedure of the annotation has been explained to all annotators beforehand.

### B.3 Label Suggestion Model

We used the `german-bert-cased` BERT base model which was pretrained on a German Wikipedia Dump (6GB), an OpenLegalData dump (2.4GB) and news articles (3.6GB). It was trained for 810k steps with a batch size of 1024 for sequence length 128 and 30k steps with sequence length 512. It outperformed the multilingual version of BERT on several downstream tasks using German data (GermEval-2018<sup>12</sup>, GermEval-2014 NER<sup>13</sup>, 10kGNAD<sup>14</sup>). More information can be found at the corresponding website<sup>15</sup>.

For our setup, we performed a random hyperparameter search using the following combinations:

- learning rate: [0.01, 0.1, 0.001, 0.0001,

<sup>12</sup><https://projects.fzai.h-da.de/iggsa/germeval-2018/>

<sup>13</sup><https://sites.google.com/site/germeval2014ner/data>

<sup>14</sup><https://tblock.github.io/10kGNAD/>

<sup>15</sup><https://deepset.ai/german-bert>

0.00001, 0.00005, 0.00008]

- batch size: [4,8,16]

### C Label Rejections

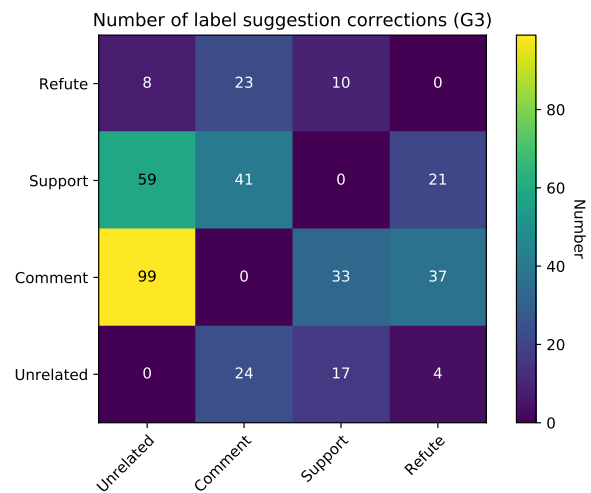


Figure 9: Number of rejected label suggestions for group G3. The x-axis displays the corrected label and the y-axis the label suggestion. For example, the upper left corner shows that 8 suggestions of label `Refute` were corrected as `Unrelated` by the users.

Figure 9 displays how student annotators from G3 corrected label suggestions, per category. As discussed in Section 5.2 we observe a similar pattern as for annotator group G2. The majority of label corrections are for the predicted category `Comment` or corrections for a wrongly predicted stance (e.g., predictions of `Support` or `Refute`).