The Impact of Stance Object Type on the Quality of Stance Detection

Maxwell A. Weinzierl, Sanda M. Harabagiu
Human Language Technology Research Institute, Department of Computer Science,
The University of Texas at Dallas
800 W. Campbell Rd., Richardson, TX 75080
{maxwell.weinzierl, sanda}@utdallas.edu

Abstract

Stance as an expression of an author's standpoint and as a means of communication has long been studied by computational linguists. Automatically identifying the stance of a subject toward an object is an active area of research in natural language processing. Significant work has employed topics and claims as the object of stance, with frames of communication becoming more recently considered as alternative objects of stance. However, little attention has been paid to finding what are the benefits and what are the drawbacks when inferring the stance of a text towards different possible stance objects. In this paper we seek to answer this question by analyzing the implied knowledge and the judgments required when deciding the stance of a text towards each stance object type. Our analysis informed experiments with models capable of inferring the stance of a text towards any of the stance object types considered, namely topics, claims, and frames of communication. Experiments clearly indicate that it is best to infer the stance of a text towards a frame of communication, rather than a claim or a topic. It is also better to infer the stance of a text towards a claim rather than a topic. Therefore we advocate that rather than continuing efforts to annotate the stance of texts towards topics, it is better to use those efforts to produce annotations towards frames of communication. These efforts will allow us to better capture the stance towards claims and topics as well.

Keywords: Stance, Framing, Social Media

1. Introduction

Stance is defined by Biber and Finegan (1988) as the expression of an author's standpoint and judgment towards a given proposition. In Bois (2007) stance is defined as *"a public act by a social actor, achieved dialogically through overt communicative means, of simultaneously evaluating objects, positioning subjects (self and others), and aligning with other subjects, with respect to any salient dimension of the sociocultural field"*, indicating that the stance-taking process is affected not only by personal opinions or attitudes, but also by other sociocultural aspects, e.g. cultural norms, social roles and social actions, etc. As pointed out in Hardalov et al. (2021), stance was defined differently when it was used in different settings for capturing the attitudes expressed in texts, e.g. the attitudes of politicians towards a newly proposed legislation (Somasundaran and Wiebe, 2010), of customers regarding new products (Somasundaran and Wiebe, 2009), or of the general public towards public health measures such as vaccinations (Glandt et al., 2021; Weinzierl et al., 2023).

Stance is used for discovering attitudes from various sources, ranging from social media, to debates or news articles, as shown in Hardalov et al. (2021). Regardless of the context in which it was used, stance always has a subject and an object. The subject of stance can be the speaker in a conversation or the author of a social media post. The stance object, as reported in Hardalov et al. (2021, 2022); Liu et al. (2023), can be a controversial entity, concept, idea, event, article headline or claim.

The different types of stance objects have led AL-Dayel and Magdy (2021) to observe that currently there are two major approaches used for stance detection: (1) sentiment-based approaches, having the purpose to uncover the implicit sentiment (favor/against) expressed in texts towards the stance object, which can be a topic or target (Mohammad et al., 2016; Sobhani et al., 2017; Allaway and McKeown, 2020; Li et al., 2021); or (2) position-based approaches, which determine if a text agrees/disagrees with a given claim or an article headline (Ferreira and Vlachos, 2016; Chen et al., 2019; Hanselowski et al., 2019; Conforti et al., 2020). These approaches exploit the distinction of semantics pertaining to stance values annotated across existing datasets. But they do not tackle the forms of judgments that capture the stance subject's standpoint towards the stance object, as indicated by the stance definition from Biber and Finegan (1988).

In this paper we take the position that, in order to follow the stance definition from Biber and Finegan (1988), for each type of stance object, different forms of judgments should be expected when deriving the stance of a text towards that object. Because Hardalov et al. (2021) have shown that the quality of stance detection is dependent on the

LREC-COLING 2024, pages 15942–15954

source of the texts (e.g. whether they originate on social media or if their provenance is from debate forums or from news articles) we shall restrict our analysis of stance detection on microblogs originating from X, formerly known as Twitter. Our decision is motivated by (a) the fact that the majority of stance-annotated datasets use this source and (b) some of the largest current annotated datasets (Conforti et al., 2020; Zheng et al., 2022) also use this source.

In addition, our position is that there should not be only two major approaches to stance detection, but three. While the first approach should consider stance objects as topics or targets; the second approach should consider stance objects represented by claims; and the third approach should use frames of communication as a new type of stance object.

Topics (or targets) as the first stance object type, typically correspond to the name of a controversial political figure, e.g. Hillary Clinton, Donald Trump, or a noun phrase, e.g. "gun control" or "feminism". The issue with topics as stance objects is that they are not always explicit in the text, and furthermore, complex judgments are involved in determining the stance towards them. For example, given the following microblog text:

> You're natural immune system is far superior than a vaccine in fighting Covid-19. Herd immunity has been around long before vaccines existed.

determining from it the stance towards the controversial topic "COVID-19 vaccines" requires first the recognition that the author of the tweet selected to address two aspects about COVID-19 vaccines: (1) compliance in recognizing the necessity of getting vaccinated; and (2) the calculation that weights the benefits of getting vaccinated vs. waiting for herd immunity to be established. For the first aspect, the text author claims that the natural immune system is superior to the COVID-19 vaccine, and therefore it is unnecessary; while for the second aspect, the text author claims that vaccines are not superior to herd immunity. Based on these claims, it becomes clear that the stance of the microblog text towards the topic of "COVID-19 vaccines" is *Reject*.

As pointed out in Küçük and Can (2021) a distinction is made between topics and position statements, e.g. "We should disband NATO", which we shall consider as claims. Govier (2013) defines a claim as a statement that is in dispute and that we are trying to support with reasons. When considering the same microblog as before, to infer the stance towards the claim:

> **CLAIM**: *There are better options than the Covid-19 vaccine.*

requires inferring the reasons surrounding this claim and how the reasons align with the statements from the microblog. Previous efforts (Hasan and Ng, 2014; Habernal et al., 2018) for inferring claim reasons and their interaction with stance have shown that this is a difficult task, with resulting F-scores in the low 50s. When considering claims as the second stance object type, the reasons that need to be inferred inform the value of the stance. Knowing that the reason for this claim is that it is better to wait for herd immunity, the microblog stance towards the claim is *Support*. Please note that the stance towards this claim participated in the judgments of the stance towards the topic "COVID-19 vaccines" attributed to the same microblog.

The third type of stance object that we consider are Frames of Communication (FoCs). As defined in Entman (1993), FoCs *"define problems – determine what a causal agent is doing in terms of costs and benefits, measured in terms of common cultural values; diagnose causes – identify the forces creating the problems; make moral judgments – evaluate causal agents and their effects; and suggest remedies – offer and justify treatments for the problems and predict their likely effects."* In the articulation of the FoCs, the problems are not explicit, but the causes, or the reasons, of the highlighted problems are always explicit (Entman, 1993). Thus FoCs provide explicitly the reasons needed to be inferred by claims. This entails that the judgments of the stance of a text towards an FoC are much simpler than the judgments required to be inferred when determining the stance towards a claim or a topic. When considering the same microblog as before, the stance towards the FoC:

> **Frame of Communication:** *Preference for getting COVID-19 and fighting it off than getting vaccinated.*

would be inferred as *Accept* because the reasons for the problems of vaccine compliance and calculation are explicit in the FoC. We note that entailment between microblogs and FoCs was considered before in the COVIDLies experiments (Hossain et al., 2020), where misconceptions gathered from Wikipedia were in fact FoCs.

In this paper we position our belief that it is easier to infer the stance of a text towards a Frame of Communication (FoC) than towards a claim, which in turn is easier to infer than towards a topic. To prove our belief, in Section 2 we detail the stance framework that relies on (a) implicit knowledge and (b) different forms of judgments when deriv-

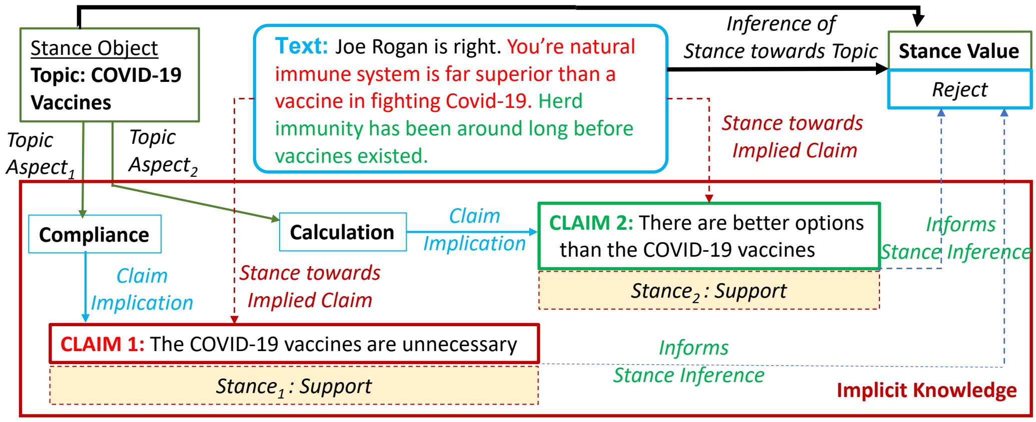

Figure 1: Inference of Stance towards Topics.

ing the stance towards each stance object type. In Section 3, we present a general stance detection method, capable of inferring the stance of a text towards each type of stance object that we have considered. In Section 4 we showcase our results and discuss lessons learned, while Section 5 summarizes the conclusions supporting our position regarding stance detection towards different stance objects.

2. Stance Object Types

2.1. Stance Object Type 1: Topics

The stance of a *Text* towards a *Topic* is inferred by assigning a *Stance Value*, reflecting acceptance, rejection or no stance at all towards the *Topic*, discussed in the *Text*, as shown in Figure 1. This inference was cast as a classification problem, starting from the seminal work reported in Mohammad et al. (2016) and gaining increasing popularity (Ghosh et al., 2019; Conforti et al., 2020; Li et al., 2021; Allaway and McKeown, 2020). The research problems of cross-topic stance detection (Hardalov et al., 2021; Conforti et al., 2020) or cross-language stance detection (Zotova et al., 2020) were derived from this formalization of stance detection. Significant research employed knowledge bases, such as Wikipedia, to assist in identifying stance values by providing additional *Topic*-specific knowledge (Bar-Haim et al., 2017; Hanawa et al., 2019; Fromm et al., 2019; He et al., 2022). However, until now, no inspection of the kind of implicit knowledge required by the inference of the stance value derived from a *Text* towards a *Topic* has been performed. Uncovering the implicit knowledge allows us to discover the *judgments* performed by human annotators when gold-standard annotations of stance values are produced.

When considering topics as stance objects, in Figure 1 we show that first, we need to take into account that for each *Topic*, one or several *Topic aspects* are addressed in the *Text*. Figure 1 shows that for the topic of COVID-19 Vaccines, the aspects of the topic addressed in the *Text* are the compliance with vaccination and the calculation regarding vaccination. This is the first form of implied knowledge that is considered. From each of these implied topic aspects, a claim is also implied from the text. Claim$_1$ is derived from the sentence highlighted in red, while Claim$_2$, is derived from the sentence highlighted in green. Figure 1 shows the value of the stance of the text towards both claims. Implied claims and their stance represent the second form of implied knowledge. Finally, deciding the stance value of the *Text* towards the *Topic* requires making judgments with all the implied knowledge. In the example illustrated in Figure 1 both claims are supported by the text, but because both these claims are against vaccination, the stance of the text towards the *Topic* is *Reject*.

2.2. Stance Object Type 2: Claims

When inferring the stance value of a *Text* towards a *Claim* we took into account the interactions between the *reasons* of the *Claim* and the *Stance Value* of the *Text* towards the *Claim*, previously advocated by Hasan and Ng (2014). In Figure 2 we show that the implicit knowledge in this case are the reason(s) supporting the *Claim*, which enable the judgment that the *Text* supports or rejects the *Claim*.

The definition of claims was introduced in the 1950s in Toulmin's influential work on argumentation (Toulmin, 2003). A claim is considered *"an assertion that deserves our attention"*. Furthermore, we consider as stance objects only claims that are argumentative. Recent work (Ma et al., 2018) has shown that it is possible to reliably discover topic-relevant claims. To complete the implied knowledge required by stance detection towards claims, we analyze, as shown in Figure 2, how claim reasons can be discovered as well as how they participate in the decision of the stance values.

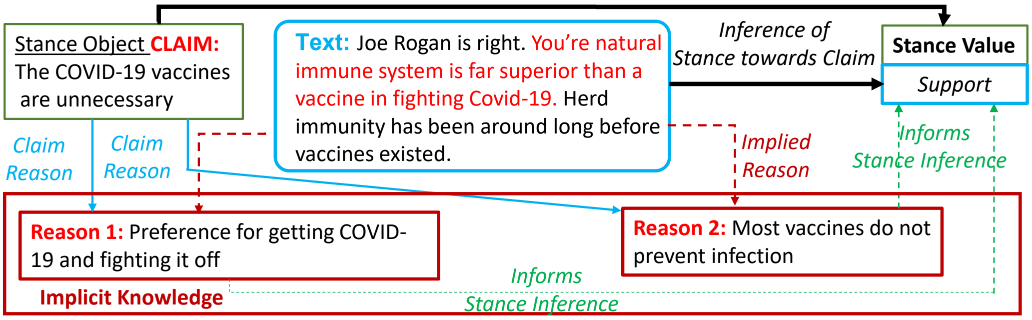

Figure 2: Inference of Stance towards Claims.

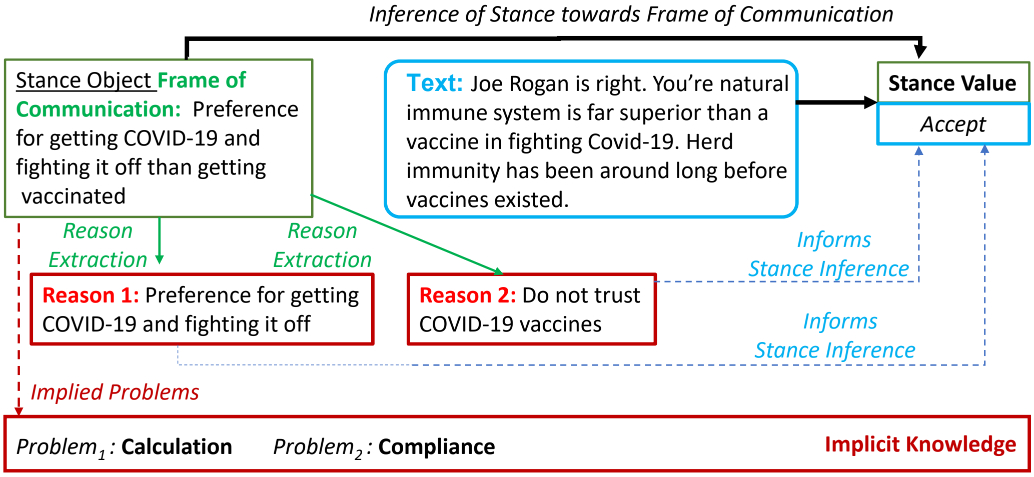

Figure 3: Inference of Stance towards Frames of Communication.

Although research addressing claim detection in texts, as part of argumentation mining, has attracted plenty of interest (Chakrabarty et al., 2019; Chen et al., 2019; Gupta et al., 2021; Khan et al., 2022; Gangi Reddy et al., 2022), the inference of claim reasons based on the textual context of a claim, has not received as much attention (Hasan and Ng, 2014). Figure 2 exemplifies how the two reasons explaining the claim can be derived from the red-highlighted sentence from the text, explaining why the stance value of the *Text* towards the *Claim* is *Support*.

An interesting new direction of performing claim verification by generating explanations (Atanasova et al., 2020) may be further extended to also generate claim reasons, which could be considered as implicit knowledge used in stance discovery. Furthermore, an extension of the discovery of the various perspectives of a claim, as introduced in Chen et al. (2019), to account for the explanations and reasons of the claim perspectives could also be beneficial for uncovering the implied knowledge illustrated in Figure 2.

2.3. Stance Object Type 3: Frames of Communication

Because Entman (1993) notes that framing involves (a) selecting salient problems and (b) articulating their causal interpretation, we expect that from each Frame of Communication - *FoC* - we can detect: (1) the implied problems; as well as (2) the reason(s) for each problem, which are explicit. Therefore, it is much simpler to infer the stance value when more of the knowledge is explicit.

Figure 3 shows how the stance value is inferred for an exemplified *FoC*. It can be seen that the first reason is derived from the text fragment *"You're natural immune system is far superior than a vaccine in fighting Covid-19"*, while the second reason can be entailed from the text fragment *"Herd immunity has been around long before vaccines existed"*. Based on these entailments of the reasons expressed in the FoC, the stance value of the *Text* towards the *FoC* is *Accept*.

When comparing the implicit knowledge and the judgments required for assigning the stance value when the stance object is (a) a topic, illustrated in Figure 1; (b) a claim, illustrated in Figure 2; or (c) a frame of communication, illustrated in Figure 3, it is evident that the easiest way of inferring the stance should be when its object is a frame of communication.

3. General Stance Detection

We call general stance detection a stance detection method capable of judging the stance of a text towards (a) a topic; (b) a claim, or (c) a Frame of Communication (FoC). A prerequisite of designing a supervised method for general stance detection is the availability of a text corpus annotated with all these three Stance Object Types (SOTs). Unfortunately, such a corpus does not exist. There are numerous corpora (Mohammad et al., 2016; Conforti et al., 2020; Li et al., 2021; Allaway and McKeown, 2020; Glandt et al., 2021) that annotates stance towards topics (or targets), some corpora (Mohtarami et al., 2019; Hanselowski et al., 2019; Mutlu et al., 2020; Chen et al., 2019) that annotated stance towards claims, and only one corpus, namely the CoVaxFrames dataset (Weinzierl and Harabagiu, 2022c), that annotates stance towards FoCs. For any of these existing stance-annotated corpora, additional annotations are needed for SOT_1=*topic*; SOT_2=*claim*; or SOT_3=*FoC*. However, we took into account the analysis detailed in Section 2 which indicates that the implicit knowledge required when judging the stance of SOT_2 can use the explicit knowledge used by SOT_3, and similarly the stance judgments used for SOT_1 can benefit from the judgments and knowledge available for SOT_2 and SOT_3. In this section we detail how this knowledge is used and how these judgments are made.

Extended annotations: Two classes of additional annotations were produced on the CoVaxFrames dataset such that we can have complete annotations that would allow us to measure the impact of SOTs on the performance of stance detection. CoVaxFrames includes 14,180 tweets with annotated stance values towards 113 FoCs, covering only one topic, namely the COVID-19 vaccine. In CoVaxFrames, the stance takes values from the set {*Accept, Reject, No Stance*}. Our first class of annotations aimed to eliminate all the implicit knowledge required when deciding the stance value of a *Text*$_i$ towards an *FoC*$_j$ by revealing the set of salient problems of each FoC available in CoVaxFrames. We believe that knowing the salient problems addressed by FoCs is important. This is because, based on the analysis presented in Section 2, we concluded that the aspects of topics used when deciding the stance of a text towards a topic correspond to the salient problems that could be selected by an FoC to interpret the text. Furthermore, for each problem (or topic aspect), as shown in Figure 1, claims are made, thus highlighting the inter-relations between the stance of a text towards a topic, a claim, or a FoC. The second class of annotations produced stance judgment towards claims and topics.

Annotating the problems of COVID-19 vaccines: The salient problems addressed by the FoCs annotated in CoVaxFrames are informed by the model of vaccine hesitancy Geiger et al. (2022), which is characterized by seven problems, or factors, that increase or decrease an individual's likelihood of getting vaccinated. The problems are: (1) *Confidence* in vaccines; (2) *Complacency* to getting vaccinated; (3) *Constraints* that make vaccination difficult or costly; (4) *Calculation*, in which personal costs and benefits of vaccination are weighted; (5) *Collective Responsibility*, that protects others from getting infected or sick; (6) *Compliance*, resulting from societal monitoring and sanctioning of people that are not vaccinated; and (7) *Conspiracy* thinking and belief in vaccine misinformation. Four researchers with expertise in vaccine communication were tasked with deciding which of the 7 problems are salient to each of the 113 FoCs available in CoVaxFrames. Researchers obtained a very high inter-annotator agreement of 81%, measured as the percentage of judged problems aligning with the majority decision, with remaining disagreements adjudicated through discussion.

Annotating stance towards additional SOTs: Aims at judging (a) the *stance(Text*$_i$,*Claim*$_j$) and (b) the *stance(Text*$_i$,*Topic*$_T$) when knowing the *stance(Text*$_i$,*FoC*$_k$). However, before producing these judgments, claims need to be generated for all text covering the *Topic*$_T$, which in CoVaxFrames is focusing on COVID-19 vaccines. For this purpose, for each of the 7 salient problems covered in the FoCs available in CoVaxFrames two separate claims were articulated, for a total of 14 claims. Given a problem P, one of the claims, $C_{pro}(P)$, is supportive of P, whereas the second one, $C_{con}(P)$, is contrary to P. For example, if P is *Complacency*, then C_{pro} is articulated as *"The COVID-19 vaccines are unnecessary"* while C_{con} is articulated as *"The COVID-19 vaccines are necessary to fight COVID-19"*. In this way, we have generated 14 claims against which stance can be judged, thus extending the annotations in CoVaxFrames to consider also SOT_2. Each generated claim is further detailed in Appendix A. Two annotators perform judgments included in:

Rule 1: : **if** *(stance(Text$_i$, FoC$_k$) is annotated in CoVaxFrames)*
and *(FoC$_k$ makes Problem$_x$ salient)*
and *(Claim$_j$ is generated for Text$_i$ and Problem$_x$)*
and *(Reason$_y$ of Problem$_x$, explicit in FoC$_k$ can be judged to support a reason of Claim$_j$)*
then *Stance_Value(Text$_i$,Claim$_j$) = Stance_Value(Text$_i$, FoC$_k$)*

In order to consider stance annotations for SOT_2, first we had to analyze which of the seven problems that were salient to all FoC are supporting the *Topic*$_T$, i.e. COVID-19 vaccines, and which problems are rejecting it. The vaccine communication experts decided that the set of supporting problems is P_{pro}={*Confidence, Compliance, Collective Responsibility, Calculation*} while the set of rejecting problems is P_{con}={*Complacency, Constraints, Conspiracy*}. Then the two annotators were asked to perform judgments included in:

Rule 2: : **if** *(stance(Text$_i$,Claim$_j$) is annotated in CoVaxFrames)*
and *(Claim$_j$ is generated for Problem$_x$)*
and *(Problem$_x$ can be judged to be an aspect of Topic$_T$)*
then if *(TYPE(Claim$_j$) = C$_{PRO}$)* **and** *(Problem$_x$ is in P$_{PRO}$)*
then *Stance_Value(Text$_i$, Topic$_T$) = Stance_Value(Text$_i$, Claim$_j$)*
else if *(TYPE(Claim$_j$) = C$_{CON}$)* **and** *(Problem$_x$ is in P$_{CON}$)*
then *Stance_Value(Text$_i$, Topic$_T$) = Stance_Value(Text$_i$, Claim$_j$)*
else *Stance_Value(Text$_i$, Topic$_T$) is opposite to Stance_Value(Text$_i$, Claim$_j$)*

The inter-annotator agreement between the two annotators when Rule 1 was considered was 96% when judging if *FoC*$_k$ supports *Claim*$_j$. Additionally, the inter-annotator agreement when Rule 2 was considered was 98% when judging whether *Problem*$_X$ is a salient aspect of *Topic*$_T$ for *Text*$_i$. Differences between annotators were adjudicated through discussions. The annotations for this dataset are are made available on GitHub[1].

Models of General Stance Detection: A straightforward approach to infer the stance value of a *Text*$_i$ towards any of the three SOTs is to consider three different models: (1) M_T, for which SOT=SOT_1, i.e. topics; (2) M_C, for which SOT=SOT_2, i.e. claims; and (1) M_F, for which SOT=SOT_3, i.e. frames of communication. A simple neural architecture capable of implementing each of these three models is presented in Figure 4. This simple architecture is inspired by other stance detection methods that regularly employ a domain-specific BERT (Devlin et al., 2019) model, fine-tuned for stance detection to directly infer stance values (Hossain et al., 2020; Barbieri et al., 2020; Weinzierl et al., 2021; Weinzierl and Harabagiu, 2022c,a,b).

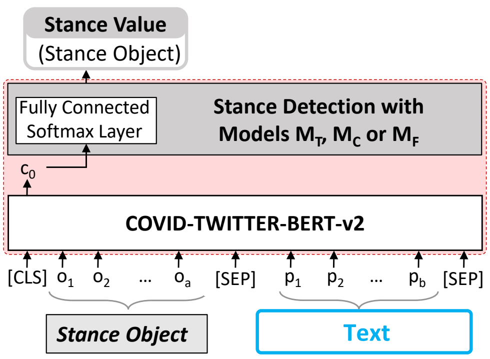

Figure 4: Direct stance detection.

[1] https://github.com/Supermaxman/co-vax-frames-claims-topics

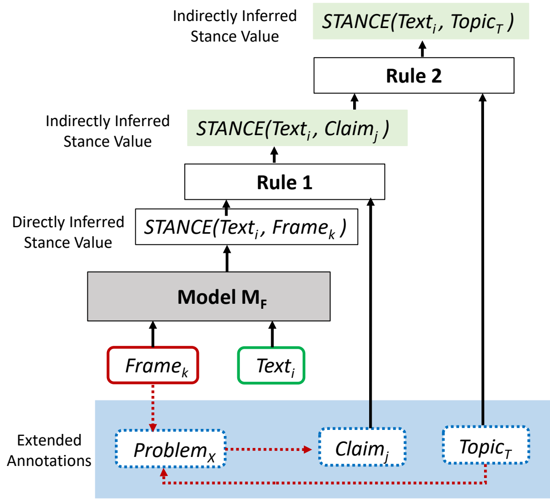

Figure 5: Indirect stance detection using M_F.

The direct detection of stance approach illustrated in Figure 4 shows that we used COVID-Twitter-BERT-v2 (Müller et al., 2023), which was further pre-trained on 97M COVID-19 tweets, providing domain-specific language modeling for tasks concerning COVID-19. In our approach, we use (a) the language describing a stance object o, e,g, a topic, a claim, or the articulation of an FoC, available from the extended annotations of CoVaxFrames; and (b) a text originating in a microblog p. Word-piece tokenization (Devlin et al., 2019) of the text of o and the text of p is provided to COVID-Twitter-BERT-v2 for generating corresponding contextualized embeddings. The contextualized embedding c_0, corresponding to the "[CLS]" token, is provided to a fully connected layer, employing a softmax activation function to predict the probabilities of the possible stance values. When the SOT(o)=SOT_1, the architecture illustrated in Figure 4 implements a model M_T, whereas when SOT(o)=SOT_2, model M_C is implemented whereas when SOT(o)=SOT_3, model M_F is implemented.

However, stance values can also be discovered indirectly through inference. As shown in Figure 5, in addition to the stance value of a *Text*$_i$ towards a *Frame*$_k$ directly inferred by model M_F, by using (a) Rule 1 or Rule 2; as well as (b) the extended annotations, we can indirectly infer two other stance values. We infer (1) the stance of the same *Text*$_i$ towards a *Claim*$_j$ generated from a *Problem*$_X$ deemed salient by *Frame*$_k$; and (2) the stance of the same *Text*$_i$ towards a *Topic*$_T$ that addresses the same aspect as *Problem*$_X$ which deemed salient by *Frame*$_k$, and from which *Claim*$_j$ was generated.

Similarly, Figure 6 illustrates how model M_C directly infers the stance of *Text*$_i$ towards a *Claim*$_j$, and indirectly it infers the stance of the same *Text*$_i$

Model Type	SOT of Inferred Stance	Macro F_1	Macro P	Macro R	Accept F_1	Accept P	Accept R	Reject F_1	Reject P	Reject R
M_F	FoCs^D	**82.85**	**81.02**	**84.78**	**90.08**	**88.56**	**91.65**	**75.62**	**73.47**	**77.90**
M_C	Claims^D	72.37	80.01	66.39	83.14	87.30	79.36	61.60	72.73	53.43
M_F	Claims^I	**84.42**	**82.48**	**86.47**	**90.40**	**89.29**	**91.54**	**78.43**	**75.67**	**81.41**
M_T	Topic^D	75.92	81.24	71.35	72.31	79.66	66.20	79.54	82.82	76.50
M_C	Topic^I	80.05	84.35	76.16	79.38	83.86	75.35	80.71	84.84	76.97
M_F	Topic^I	**87.66**	**86.38**	**89.01**	**86.16**	**83.87**	**88.59**	**89.16**	**88.90**	**89.42**

Table 1: Stance detection results when considering different Stance Object Types (SOTs).

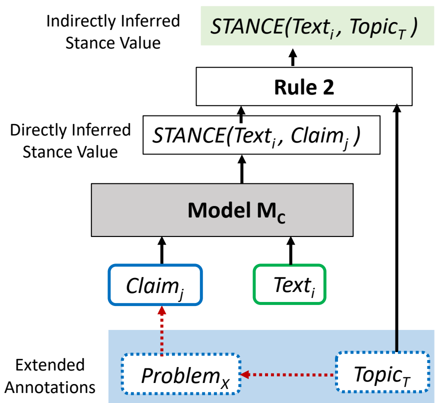

Figure 6: Indirect stance detection using M_C.

towards a *Topic*$_T$, which is addressing in *Text*$_i$ an aspect that is the same as *Problem*$_X$, which enables the generation of *Claim*$_j$

All hyperparameters of M_F, M_C and M_T were determined by maximizing stance detection performance on the development collection of the extended annotations of CoVaxFrames, detailed in Appendix B.

4. Experimental Results and Discussion

Experimental Results: To measure the impact of SOTs on the quality of stance detection, we trained each system on the training collection from CoVaxFrames and evaluated on the test collection from CoVaxFrames, enriched with the extended annotations described in Section 3. We measured Precision (P), Recall (R), and F_1 metrics for detecting the *Accept* and *Reject* values of stance. We also compute a Macro-averaged Precision, Recall, and F_1 score for stance detection in general. The results are listed in Table 1, where the bolded numbers represent the best results obtained for each SOT and each model detailed in Section 3. The Table shows when the stance value was determined directly, marked by the D superscript, or indirectly, marked by the I superscript.

The results from Table 1 empirically demonstrate the benefits of considering FoCs as the ideal SOT towards which stance detection is performed, because the model M_F outperforms all other models. Also model M_C is superior to model M_T, given their performance. The results confirm our intuitions that stance is more difficult to detect when more implicit knowledge is required.

Results for stance detection towards the *"COVID-19 Vaccines"* topic also show how indirect stance detection methods based on the M_C or M_F models outperform direct methods using the M_T model. When inferring the stance of a text towards the *"COVID-19 Vaccines"* topic, if model M_T is used to directly infer the stance, its results, measured by the Macro F_1-score, are worse by almost 5 points than when the stance is indirectly inferred with model M_C. Also these results are much worse than when the stance is indirectly inferred by model M_F, yielding a difference of almost 12 points in Macro F_1-score. All these differences quantify how much performance is lost when the implicit knowledge used for stance detection is hard to capture.

When detecting the stance of a text towards a claim, the model M_F improves by more than 12 points of F_1-score the results of direct stance detection over the results of the model M_C which infers directly the stance.This difference in performance quantifies the difficulty of inferring how the claim reasons interact with the claim stance value - a task known to be challenging, as reported in (Hasan and Ng, 2014).

Discussion: The findings of our experiments explain why current stance detection methods struggle to achieve F_1-scores above ~ 76, when considering the SOT to be topics. For example, the best stance detection system (Loureiro et al., 2022) operating on the SemEval-16 Task-6 (Mohammad et al., 2016) dataset achieved an F_1 score of 72.9 by incorporating continuous learning into language models operating on social media, while Liu et al. (2023) reported an F_1 score of 71.1 with significant effort involved in enriching the training dataset. Even the scaling-up of topic-based stance detection datasets by annotating additional examples does not appear to resolve this issue. The Will-They-Won't-They (WT-WT)

(Conforti et al., 2020) dataset, which consists of 51,284 microblogs annotated with stance values towards five merger and acquisition topics, measured an upper-bound of human performance as 75.2 F_1-score based on human-level agreement, with stance detection achieving an F_1 score of 74.9 by incorporating stock market signals (Conforti et al., 2022). Furthermore, stance detection models operating on claims from Snopes achieved an F_1 score of 75.07 (Hardalov et al., 2021); close to the human-level performance of an F_1 score of 80.2 (Hanselowski et al., 2019).

□ We therefore advocate to annotate text datasets with stance values by considering the stance objects to be Frames of Communication (FoCs) rather than topics. For this purpose, we also note that there is a need to develop automatic methods capable of recognizing the FoCs evoked in text, which would greatly facilitate the annotations. For example, recently Weinzierl and Harabagiu (2024) utilized GPT-4 (OpenAI, 2023), a Large Language Model (LLM), to fully discover and articulate FoCs on CoVaxFrames. This LLM-based approach employed Chain-of-Thought (CoT) prompting (Wei et al., 2022) with In-Context Active Curriculum Learning (CoT-ICACL), demonstrating human-level automatic frame discovery and articulation with a thorough manual evaluation of the 292 discovered FoCs.

SOT impact on inter-annotator agreement: We note that the inter-agreement annotations when FoCs are considered as the SOT is superior to the stance annotations performed when the SOTs was a topic. For example, Weinzierl and Harabagiu (2022c) reports the inter-annotator Cohen's Kappa score when judging stance towards FoCs as 0.67. But for the SemEval-16 Task-6 dataset which annotates stance toward topics (Mohammad et al., 2016) reports an agreement accuracy of 73.1, corresponding to an estimated Cohen's Kappa score of 0.58. For the annotations performed on the P-Stance dataset (Li et al., 2021), where the SOT was also the topic, the inter-annotator agreement was reported as having a Krippendorff's alpha of 0.6. Inter-annotator agreement for stance of text towards claims, as reported for the Snopes dataset (Hanselowski et al., 2019), had Cohen's Kappa of 0.7, therefore better than the agreement obtained when the SOT is a topic.

□ It appears that human annotators have more difficulty inferring stance values of texts towards topics than towards FoCs or claims. This difficulty appears to be translated in the performance of stance detection systems judging the stance of texts towards these different SOTs.

Complexity of Implied Knowledge: For all SOTs stance detection requires implied knowledge. However, sometimes, the required implied knowledge is very complex. We exemplify this with a microblog from the SemEval-16 Task-6 dataset:

we remind ourselves that love means to be willing to give until it hurts - Mother Teresa

This microblog is annotated as holding an *Against* stance towards the topic of *Abortion*. Recognizing this stance value requires significant background knowledge: First, involving an implicit aspect surrounding abortion concerning suffering for loved ones, and furthermore recognizing that this microblog contains a reason in support of this aspect. This is a direct quote from Mother Teresa in a famous speech concerning abortion at a National Prayer Breakfast in Washington, D.C., on February 3rd, 1994[2]. Similarly, the following microblog holds an *Against* stance towards the topic of *Hillary Clinton*:

@user @user Never forget @user #Benghazi Shackle the criminal element

Recognizing this microblog is *Against* the topic of Hillary Clinton requires understanding the implicit aspect of controversies implicating Hillary Clinton. Furthermore, one must know that Benghazi was a tragic event for which Hillary Clinton was blamed, and that this microblog is calling for others to *"Never forget"* this event - a likely reference to the common saying concerning the tragic events of 9/11 (Fink and Mathias, 2002).

□ The role of the knowledge implied by a text, a topic, a claim or a frame of communications is central to the inference of the stance of a text towards any of the SOTs we have considered. We advocate for the need of additional research that can reveal the implied knowledge while reasoning with it to establish the stance. It is also important to develop capabilities to infer different forms of implied knowledge, including social bias (Sap et al., 2020).

Limitations: Although our experiments have shown that FoCs are the most beneficial SOT when inferring stance, we have several experimental limitations. The first limitation stems from the usage of a single dataset, namely CoVaxFrames. This creates a second limitation, because the CoVaxFrames dataset refers only to the single topic of *"COVID-19 Vaccines"*. Therefore, we could not demonstrate that our findings operate across a variety of topics, like the VAST dataset (Allaway and McKeown, 2020) does, as in its 23,525 news article comments annotated with stance values, it addresses 5,634 topics.

A third limitation derives from the fact that articulating FoCs, and to a lesser extent claims, can add

[2] https://www.c-span.org/video/?54274-1/national-prayer-breakfast

significant effort to the process of stance annotation when compared to topics. However, we argue that annotators must also resolve implicit aspects surrounding a topic, and furthermore must implicitly recognize reasons that support or reject claims corresponding to each aspect in order to perform accurate stance annotations towards topics. As we have discussed in Section 2, resolving these implicit relationships requires that annotators possess significant knowledge surrounding a topic. Furthermore, due to the difficulty of articulating FoCs, human annotators can make mistakes. For example, the FoC *"It takes courage both to vaccinate against COVID-19 and to refuse the vaccine"* from CoVaxFrames clearly evokes the problem of *Calculation*, but does not clearly articulate a reason in support or contrary to *Calculation*. We believe that this last limitation can be resolved by research focusing on automatically discovering and articulating FoCs from texts, similar to how current research works towards discovering claims from texts (Chakrabarty et al., 2019; Chen et al., 2019; Gupta et al., 2021; Khan et al., 2022; Gangi Reddy et al., 2022). This new line of research would enable the stance detection systems to reach near-human levels of performance, as our experiments indicate. Towards this end, recently the first methodology for automatically discovering and articulating FoCs was proposed by Weinzierl and Harabagiu (2024), which utilized LLMs to discover and articulate FoCs at human-level performance.

5. Conclusion

In this paper, we have presented the position that objects of stance have a significant impact on the quality of stance detection. We have advocated the need to consider the stance of a text towards a frame of communication rather than towards a claim or a topic. Our position is backed by experimental results that clearly indicate that it is easier to infer the stance of a text towards a frame of communication than towards a claim or a topic. We hope that the lessons learned from our analysis and the experimental results will provide important insights to researchers interested in learning to infer stance from texts and will lead to new research directions that consider the interaction between stance and topics, claims, and frames of communication, which have an important role in the interpretation of texts.

6. Bibliographical References

Abeer ALDayel and Walid Magdy. 2021. Stance detection on social media: State of the art and trends. *Information Processing & Management*, 58(4):102597.

Emily Allaway and Kathleen McKeown. 2020. Zero-Shot Stance Detection: A Dataset and Model using Generalized Topic Representations. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8913–8931, Online. Association for Computational Linguistics.

Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. Generating fact checking explanations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7352–7364, Online. Association for Computational Linguistics.

Roy Bar-Haim, Lilach Edelstein, Charles Jochim, and Noam Slonim. 2017. Improving claim stance classification with lexical knowledge expansion and context utilization. In *Proceedings of the 4th Workshop on Argument Mining*, pages 32–38, Copenhagen, Denmark. Association for Computational Linguistics.

Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1644–1650, Online. Association for Computational Linguistics.

Douglas Biber and Edward Finegan. 1988. Adverbial stance types in english. *Discourse Processes*, 11(1):1–34.

John W. Du Bois. 2007. The stance triangle. *Pragmatics and beyond. New series*, 164:139–182.

Tuhin Chakrabarty, Christopher Hidey, and Kathy McKeown. 2019. IMHO fine-tuning improves claim detection. 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 558–563, Minneapolis, Minnesota. Association for Computational Linguistics.

Sihao Chen, Daniel Khashabi, Wenpeng Yin, Chris Callison-Burch, and Dan Roth. 2019. Seeing things from a different angle:discovering diverse perspectives about claims. 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 542–557, Minneapolis, Minnesota. Association for Computational Linguistics.

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.

Costanza Conforti, Jakob Berndt, Mohammad Taher Pilehvar, Chryssi Giannitsarou, Flavio Toxvaerd, and Nigel Collier. 2022. Incorporating stock market signals for Twitter stance detection. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4074–4091, Dublin, Ireland. 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.

Robert M. Entman. 1993. Framing: Toward clarification of a fractured paradigm. *Journal of Communication*, 43(4):51–58.

William Ferreira and Andreas Vlachos. 2016. Emergent: a novel data-set for stance classification. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1163–1168, San Diego, California. Association for Computational Linguistics.

Mitchell. Fink and Lois. Mathias. 2002. *Never forget : an oral history of September 11, 2001*. Regan Books, New York.

Michael Fromm, Evgeniy Faerman, and Thomas Seidl. 2019. Tacam: Topic and context aware argument mining. In *IEEE/WIC/ACM International Conference on Web Intelligence*, WI '19, page 99–106, New York, NY, USA. Association for Computing Machinery.

Revanth Gangi Reddy, Sai Chetan Chinthakindi, Zhenhailong Wang, Yi Fung, Kathryn Conger, Ahmed ELsayed, Martha Palmer, Preslav Nakov, Eduard Hovy, Kevin Small, and Heng Ji. 2022. NewsClaims: A new benchmark for claim detection from news with attribute knowledge. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6002–6018, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Mattis Geiger, Franziska Rees, Lau Lilleholt, Ana P. Santana, Ingo Zettler, Oliver Wilhelm, Cornelia Betsch, and Robert Böhm. 2022. Measuring the 7cs of vaccination readiness. *European Journal of Psychological Assessment*, 38(4):261–269.

Shalmoli Ghosh, Prajwal Singhania, Siddharth Singh, Koustav Rudra, and Saptarshi Ghosh. 2019. Stance detection in web and social media: A comparative study. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction: 10th International Conference of the CLEF Association, CLEF 2019, Lugano, Switzerland, September 9–12, 2019, Proceedings*, page 75–87, Berlin, Heidelberg. Springer-Verlag.

Kyle Glandt, Sarthak Khanal, Yingjie Li, Doina Caragea, and Cornelia Caragea. 2021. Stance detection in COVID-19 tweets. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1596–1611, Online. Association for Computational Linguistics.

Trudy Govier. 2013. *A practical study of argument, enhanced edition*, 7 edition. Wadsworth Publishing, Belmont, CA.

Shreya Gupta, Parantak Singh, Megha Sundriyal, Md. Shad Akhtar, and Tanmoy Chakraborty. 2021. LESA: Linguistic encapsulation and semantic amalgamation based generalised claim detection from online content. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3178–3188, Online. Association for Computational Linguistics.

Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. The argument reasoning comprehension task: Identification and reconstruction of implicit warrants. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1930–1940, New Orleans, Louisiana. Association for Computational Linguistics.

Kazuaki Hanawa, Akira Sasaki, Naoaki Okazaki, and Kentaro Inui. 2019. Stance detection attending external knowledge from wikipedia. *Journal of Information Processing*, 27:499–506.

Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. A richly annotated corpus for different tasks in automated fact-checking. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 493–503, Hong Kong, China. Association for Computational Linguistics.

Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2021. Cross-domain label-adaptive stance detection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 9011–9028, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Momchil Hardalov, Arnav Arora, Preslav Nakov, and Isabelle Augenstein. 2022. A survey on stance detection for mis- and disinformation identification. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1259–1277, Seattle, United States. Association for Computational Linguistics.

Kazi Saidul Hasan and Vincent Ng. 2014. Why are you taking this stance? identifying and classifying reasons in ideological debates. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 751–762, Doha, Qatar. Association for Computational Linguistics.

Zihao He, Negar Mokhberian, and Kristina Lerman. 2022. Infusing knowledge from Wikipedia to enhance stance detection. In *Proceedings of the 12th Workshop on Computational Approaches to Subjectivity, Sentiment & Social Media Analysis*, pages 71–77, Dublin, Ireland. Association for Computational Linguistics.

Tamanna Hossain, Robert L. Logan IV, Arjuna Ugarte, Yoshitomo Matsubara, Sean Young, and Sameer Singh. 2020. COVIDLies: Detecting COVID-19 misinformation on social media. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.

Kashif Khan, Ruizhe Wang, and Pascal Poupart. 2022. WatClaimCheck: A new dataset for claim entailment and inference. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1293–1304, Dublin, Ireland. Association for Computational Linguistics.

Dilek Küçük and Fazli Can. 2021. Stance detection: Concepts, approaches, resources, and outstanding issues. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, page 2673–2676, New York, NY, USA. Association for Computing Machinery.

Yingjie Li, Tiberiu Sosea, Aditya Sawant, Ajith Jayaraman Nair, Diana Inkpen, and Cornelia Caragea. 2021. P-stance: A large dataset for stance detection in political domain. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2355–2365, Online. Association for Computational Linguistics.

Zhengyuan Liu, Yong Keong Yap, Hai Leong Chieu, and Nancy Chen. 2023. Guiding computational stance detection with expanded stance triangle framework. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3987–4001, Toronto, Canada. Association for Computational Linguistics.

Daniel Loureiro, Francesco Barbieri, Leonardo Neves, Luis Espinosa Anke, and Jose Camacho-collados. 2022. TimeLMs: Diachronic language models from Twitter. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 251–260, Dublin, Ireland. Association for Computational Linguistics.

Wenjia Ma, WenHan Chao, Zhunchen Luo, and Xin Jiang. 2018. CRST: a claim retrieval system in Twitter. In *Proceedings of the 27th International Conference on Computational Linguistics: System Demonstrations*, pages 43–47, Santa Fe, New Mexico. Association for Computational Linguistics.

Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. SemEval-2016 task 6: Detecting stance in tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 31–41, San Diego, California. Association for Computational Linguistics.

Mitra Mohtarami, James Glass, and Preslav Nakov. 2019. Contrastive language adaptation for cross-lingual stance detection. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4442–4452, Hong Kong, China. Association for Computational Linguistics.

Ece C. Mutlu, Toktam Oghaz, Jasser Jasser, Ege Tutunculer, Amirarsalan Rajabi, Aida Tayebi, Ozlem Ozmen, and Ivan Garibay. 2020. A stance data set on polarized conversations on twitter about the efficacy of hydroxychloroquine as a treatment for covid-19. *Data in Brief*, 33:106401.

Martin Müller, Marcel Salathé, and Per E. Kummervold. 2023. Covid-twitter-bert: A natural language processing model to analyse covid-19 content on twitter. *Frontiers in Artificial Intelligence*, 6.

OpenAI. 2023. Gpt-4 technical report.

Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. Social bias frames: Reasoning about social and power implications of language. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5477–5490, Online. Association for Computational Linguistics.

Parinaz Sobhani, Diana Inkpen, and Xiaodan Zhu. 2017. A dataset for multi-target stance detection. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 551–557, Valencia, Spain. Association for Computational Linguistics.

Swapna Somasundaran and Janyce Wiebe. 2009. Recognizing stances in online debates. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 226–234, Suntec, Singapore. Association for Computational Linguistics.

Swapna Somasundaran and Janyce Wiebe. 2010. Recognizing stances in ideological on-line debates. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 116–124, Los Angeles, CA. Association for Computational Linguistics.

Stephen E. Toulmin. 2003. *The Uses of Argument*, 2 edition. Cambridge University Press.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.

Maxwell Weinzierl and Sanda Harabagiu. 2022a. Identifying the adoption or rejection of misinformation targeting covid-19 vaccines in twitter discourse. In *Proceedings of the ACM Web Conference 2022*, WWW '22, page 3196–3205, New York, NY, USA. Association for Computing Machinery.

Maxwell Weinzierl and Sanda Harabagiu. 2022b. VaccineLies: A natural language resource for learning to recognize misinformation about the COVID-19 and HPV vaccines. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6967–6975, Marseille, France. European Language Resources Association.

Maxwell Weinzierl and Sanda Harabagiu. 2024. Discovering and articulating frames of communication from social media using chain-of-thought reasoning. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1617–1631, St. Julian's, Malta. Association for Computational Linguistics.

Maxwell Weinzierl, Suellen Hopfer, and Sanda M. Harabagiu. 2021. Misinformation adoption or rejection in the era of covid-19. *Proceedings of the International AAAI Conference on Web and Social Media*, 15(1):787–795.

Maxwell A. Weinzierl and Sanda M. Harabagiu. 2022c. From hesitancy framings to vaccine hesitancy profiles: A journey of stance, ontological commitments and moral foundations. *Proceedings of the International AAAI Conference on Web and Social Media*, 16(1):1087–1097.

Maxwell A. Weinzierl, Suellen Hopfer, and Sanda M. Harabagiu. 2023. Scaling up the discovery of hesitancy profiles by identifying the framing of beliefs towards vaccine confidence in twitter discourse. *Journal of Behavioral Medicine*, 46(1):253–275.

Jonathan Zheng, Ashutosh Baheti, Tarek Naous, Wei Xu, and Alan Ritter. 2022. Stanceosaurus: Classifying stance towards multicultural misinformation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2132–2151, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Elena Zotova, Rodrigo Agerri, Manuel Nuñez, and German Rigau. 2020. Multilingual stance detection in tweets: The Catalonia independence corpus. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1368–1375, Marseille, France. European Language Resources Association.

Aspect	Generated Claims	Example FoCs
Confidence	☐ *Pro* - The COVID-19 vaccines are safe and efficient ☐ *Con* - The COVID-19 vaccine is dangerous and inefficient	☐ Scientists have been working on Coronavirus vaccines for decades. ☐ The COVID vaccine renders pregnancies risky and unsafe for unborn babies.
Complacency	☐ *Pro* - The COVID-19 vaccines are unnecessary ☐ *Con* - The COVID-19 vaccines are necessary to fight COVID-19	☐ Preference for getting COVID-19 and naturally fighting it off with your immune system than vaccinating. ☐ Given the risks of COVID-19, it is unlikely that building natural immunity is a good idea.
Constraints	☐ *Pro* - It is hard to get the COVID-19 vaccines ☐ *Con* - It is easy to get the COVID-19 vaccines	☐ Patents for the COVID-19 vaccines should be waived, but this is opposed by the pharmaceutical companies who want to profit.
Calculation	☐ *Pro* - The COVID-19 vaccines are better than other options ☐ *Con* - Other options are better than the COVID-19 vaccines	☐ Even if the vaccine was not tested for a long time, it is not worth having the lingering effects of COVID-19. ☐ Wait one year to see if there are no long-lasting side effects of the vaccine
Collective Responsibility	☐ *Pro* - We are all responsible for getting the COVID-19 vaccines ☐ *Con* - Getting the COVID-19 vaccine is a personal choice	☐ Vaccination is key in protecting yourself and others against COVID-19.
Compliance	☐ *Pro* - The COVID-19 vaccines should be required ☐ *Con* - The COVID-19 vaccines should be optional	☐ Vaccination against COVID-19 should be mandatory / compulsory. ☐ Vaccine exemptions should be available because the COVID-19 vaccines are experimental.
Conspiracy	☐ *Pro* - The COVID-19 vaccines are a conspiracy ☐ *Con* - The COVID-19 vaccines are not a conspiracy	☐ COVID-19 vaccines make you 5G compatible and Bluetooth-enabled.

Table 2: Aspects of COVID-19 vaccination from CoVaxFrames along with generated claims and examples of corresponding Frames of Communication.

7. Language Resource References

A. Generated Claims

Table 2 contains all generated claims for each aspect from the 7C model of vaccine hesitancy. Examples of FoCs that evoke each of these aspects are also included, along with whether those FoCs and claims support or are contrary to those aspects of COVID-19 vaccination. Two vaccine communication experts worked together to articulate these 14 claims to ensure they properly represent both perspectives towards each aspect of COVID-19 vaccination.

B. Hyperparameter Selection

Hyperparameters for each of the three trained stance detection models M_F, M_C, and M_T were selected by optimizing validation Macro F_1 score on CoVaxFrames. A grid search was performed over three hyperparameters: learning rate, batch size, and number of epochs trained. Learning rate varied over the following values, which were selected based on standard BERT values utilized in prior work: 5e-7, 1e-6, 5e-6, 1e-5, 5e-5, 1e-4, 5e-4. Similarly, batch size varied over 4, 8, 16, and 32, while max epochs varied over 3, 5, 10, 15, 20, and 25. Each experiment required approximately 10 minutes to train on an Nvidia Titan V GPU, resulting in around 4 days of (non-contiguous) training. The final hyperparameters are provided for each model in Table 3.

Model	Learning Rate	Batch Size	Max Epochs
M_F	1e-5	8	5
M_C	1e-4	16	20
M_T	1e-6	8	5

Table 3: Hyperparameters selected for each stance detection model.