Emotional Framing in the Spreading of False and True Claims

Akram Sadat Hosseini
Institute for Parallel and Distributed Systems
Stuttgart, Germany
akram.hosseini@ipvs.uni-stuttgart.de

Steffen Staab
Institute for Parallel and Distributed Systems
Stuttgart, Germany
University of Southampton
Southampton, UK
steffen.staab@ipvs.uni-stuttgart.de

ABSTRACT

The explosive growth of online misinformation, such as false claims, has affected the social behavior of online users. In order to be persuasive and mislead the audience, false claims are made to trigger emotions in their audience. This paper contributes to understanding how misinformation in social media is shaped by investigating the emotional framing that authors of the claims try to create for their audience. We investigate how, firstly, the existence of emotional framing in the claims depends on the topic and credibility of the claims. Secondly, we explore how emotionally framed content triggers emotional response posts by social media users, and how emotions expressed in claims and corresponding users' response posts affect their sharing behavior on social media. Analysis of four data sets covering different topics (politics, health, Syrian war, and COVID-19) reveals that authors shape their claims depending on the topic area to pass targeted emotions to their audience. By analysing responses to claims, we show that the credibility of the claim influences the distribution of emotions that the claim incites in its audience. Moreover, our analysis shows that emotions expressed in the claims are repeated in the users' responses. Finally, the analysis of users' sharing behavior shows that negative emotional framing such as anger, fear, and sadness of false claims leads to more interaction among users than positive emotions. This analysis also reveals that in the claims that trigger happy responses, true claims result in more sharing compared to false claims.

CCS CONCEPTS

• **Information systems** → **Data analytics**; • **Computing methodologies** → **Natural language processing**.

KEYWORDS

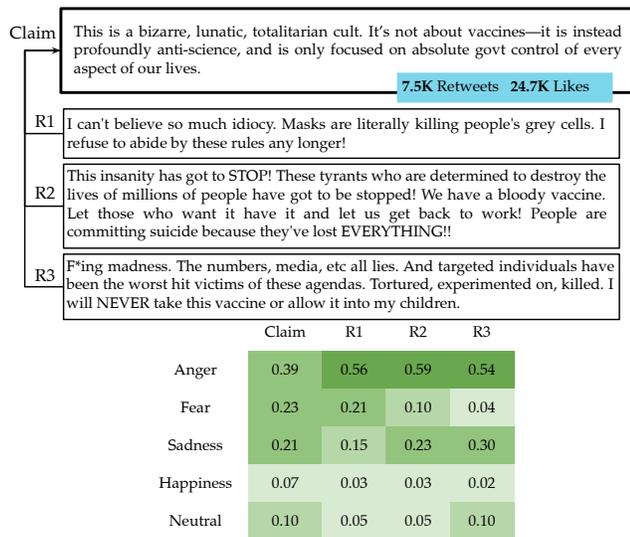Misinformation, Emotion Analysis, Social Media

1 INTRODUCTION

The sheer volume of false claims published on the Internet has affected the social behavior of online users by manipulating public sentiment [24], influencing people's perceptions, and distorting their awareness and decision-making [55]. False claims[1] comprise intentionally false statements of facts or knowingly fabricated statements to deceive users into believing they are true [26]. But what makes false claims so powerful?

Let us consider the following tweet by Senator Ted Cruz, criticizing COVID-19 guidance urging vaccinated people to continue wearing masks, saying: *"This is a bizarre, lunatic, totalitarian cult. It's not about vaccines—it is instead profoundly anti-science, and is only focused on absolute govt control of every aspect of our lives."* This statement, widely shared in December 2020, has been rated as false in PolitiFact (www.politifact.com), a reputed US political fact-checking website. This claim tries to inspire strong emotions such as anger, fear, and sadness in the audience through emotionally loaded words like 'bizarre', 'lunatic', and 'totalitarian'.

Communication literature has shown that reading an emotionally loaded message makes readers rely on the emotions in the message to form their attitude [27]. It seems that some power of the above statement is drawn from an `emotional framing` that sways its audience to accept the tweet as being true, to like and share the tweet, and to respond to it with likewise emotions (cf. Fig. 1). Emotional framing is a way of presenting the content of a claim by using emotional words and sentences that influence the thinking or behavior of a claim's consumers, in particular, to direct and control discourse and attach a certain emotion to the audiences' ideas and opinions about a given claim [11]. Therefore, exploring emotional framing is key to understanding the impact of misinformation on the audience.

Research on false claims has been performed within two broad areas, namely, technical and behavioral [47]. Technical research focuses on detecting the falsity of the information by using Natural Language Processing models and deceptive cues in the language of claims [52, 55], while behavioral research focuses on human aspects like belief, attitude, intentions, etc to understand the impact of false claims on the formation of opinions [47]. Our focus is on understanding how misinformation is shaped by exploring the impact of emotionally framed claims on their audience's behavior in social media. Therefore, we relate this work mostly to behavioral research.

---

[1]In the current political climate the "fake news" term has been co-opted by politicians as a strategy for labeling news sources that do not support their positions [14]. Therefore, in this paper, we adopt the term "false claims" instead of using the term fake news.

	Claim	R1	R2	R3
Claim	This is a bizarre, lunatic, totalitarian cult. It's not about vaccines—it is instead profoundly anti-science, and is only focused on absolute govt control of every aspect of our lives. **7.5K** Retweets **24.7K** Likes			
R1	I can't believe so much idiocy. Masks are literally killing people's grey cells. I refuse to abide by these rules any longer!			
R2	This insanity has got to STOP! These tyrants who are determined to destroy the lives of millions of people have got to be stopped! We have a bloody vaccine. Let those who want it have it and let us get back to work! People are committing suicide because they've lost EVERYTHING!!			
R3	F*ing madness. The numbers, media, etc all lies. And targeted individuals have been the worst hit victims of these agendas. Tortured, experimented on, killed. I will NEVER take this vaccine or allow it into my children.			

	Claim	R1	R2	R3
Anger	0.39	0.56	0.59	0.54
Fear	0.23	0.21	0.10	0.04
Sadness	0.21	0.15	0.23	0.30
Happiness	0.07	0.03	0.03	0.02
Neutral	0.10	0.05	0.05	0.10

Figure 1: An example of COVID-19 false claim and corresponding replies in Twitter with their emotional degrees.

Studies investigating the behavior of social media users with regard to misinformation showed that false claims are attention-grabbing [36] and designed to make it hard for humans to identify them as false by exploiting people's cognitive limitations, emotions, and ideological biases [42]. Using Elaboration Likelihood Model (ELM) [36] theorizes two routes through which true and false content generators persuade readers: a central route of high cognitive effort, and a peripheral route of low cognitive effort. Their results indicate that false claims favor the peripheral route by providing less information and crafting information in a way that requires less cognitive effort from the recipient. In a recent study, [50] revealed that falsehood diffused significantly further than the truth. They found that false claims inspire replies expressing intense emotions such as surprise and disgust. In a similar study, [31] found a strong correlation between experiencing heightened emotions and an increased probability of belief in false news. Furthermore, [15, 37, 46] posit that a negative tone was attractive to readers and therefore the authors of false claims emphasized strong negative emotions for persuading and misleading the audience to believe in falsified information and spurring them to act.

While these studies highlight **that** emotions in false claims are important factors that impact people's willingness to pass on information on social media, they fail to explain **how** emotions make false claims powerful such that they spread more widely than true claims. To answer this overall question, (i) we try to reproduce previous findings about negative emotions dominating false claims and try to generalize such findings over various topics. Then we turn to the effects that emotional false claims have on their audience, (ii), we study the audience's attitude that false claims evoke by investigating response posts for uttered emotions. Finally, (iii), we link these attitudes to the audience's behavior in terms of retweeting and liking claims, such that we better understand how emotions in false claims contribute to their broad spreading.

Corresponding to (i), (ii), and (iii), we formalize the following research questions:

- **RQ1:** In what way does the topic and credibility of a claim affect the emotional language used in a claim?
- **RQ2:** How do emotions expressed in users' response posts relate to emotions in false and true claims?
- **RQ3:** How and to what extent do the emotions expressed in reply to falsehoods inspire activities such as retweeting and liking a claim?

To operationalize the discovery of emotions in utterances to the fullest possible extent, we develop a computational framework, based on DepecheMood ++ [2], NRC Affect [34], and NRC Valence [33] lexicons, for measuring emotions in four emotional categories (anger, fear, sadness, and happiness). Our method unifies these lexical resources which are based on different emotion representations in a joint representation of four emotion labels, i.e. anger, fear, sadness, and happiness (more details in Section 4).

Our analysis of four real-world data sets covering different topics (politics, health, Syrian war, and COVID-19) showcases the interdependency between the topic of the claim and the emotions expressed in the claims, as well as, the influence of the claims' credibility on the strength of emotions expressed in the claims, rebuking previous findings that false claims would always focus on negative emotions. Analyzing users' emotional responses on social media to claims reveals that emotions expressed in the claims are repeated in the users' responses. Moreover, the results indicate the reciprocal effect of the claims' credibility on users' sharing behavior. Our analysis reveals negative emotions such as anger, fear, and sadness at false claims lead to more interaction among users than positive emotions. However, interestingly, in the claims expressing happy emotions, true claims triggered a higher level of interaction between users than false claims.

The rest of the paper is organized as follows: First, an overview of the related work is presented. Next, the proposed research model and hypotheses are described. Then, the research methodology including the implementation details is presented. Afterward, we discuss data analysis and results, and finally, we draw conclusions.

2 LITERATURE REVIEW

This section describes relevant research on misinformation that has been grouped into three broad areas, namely, content-, social-, and behavioral-level analysis.

2.1 Content Level Analysis

Early research on misinformation has been focused on analyzing writing styles and content of news from a linguistic perspective to develop better detection methods [47]. False claims mimic legitimate claims by imitating language styles of true claims or expressing opinions with a tone frequently used in true claims [44]. However, "language leakage" or so-called "deception cues" occur with certain verbal aspects in the content of a claim that is hard to monitor [14]. For example, a fine-grained analysis of word usage in deceptive texts by [28] revealed that deceivers are likely to use more other-oriented pronouns (you, your), instead of self-oriented pronouns (I, me), indicative of the speaker's discomfort in identifying themselves with the lying statements. In a similar study, [32, 35]

indicated that negative emotion verbs (e.g., hate, worthless, envy) and words related to "certainty" are dominant in deceptive texts, whereas, true claims contain significantly fewer hedges, subjective terms, and harmful words than false claims [49]. In a recent study, [30] demonstrated that combining different emotional features extracted from news content with word-embedding models, such as BERT, can enhance the accuracy of classifying misinformation.

2.2 Social Level Analysis

Social contextual features have been used in various linguistic-based approaches to enhance veracity detection [45]. Using a stance-based model, [22] defined a scheme to categorize types of reactions expressed toward false and true claims to infer news veracity. In a similar study, [13] showed that posts linked to deceptive sources receive more questions, appreciation, and answers from users, while posts linked to trusted sources have higher rates of elaboration, agreement, and disagreement. Furthermore, social psychology research indicates that emotion-eliciting content influences the sentiment of readers [48, 51]. Consequently, [54, 56] emphasized the importance of considering not only the emotions expressed in the content of news but also the emotions elicited in social media users when faced with the news, for the purpose of detecting its veracity. Moreover, according to [56], there are more instances of false news articles in which both the claim and the user response posts express anger, compared to real news. Arguing that users' interaction with information depends on their interest in its topic, [38] showed users express different emotions towards 10 different Twitter topics. In a similar study, [9] reports that while misinformation around scientific news reaches a higher level of diffusion faster, it also decays faster.

2.3 Behavioral Level Analysis

Increasingly, false news generators successfully persuade recipients to change their attitudinal responses in the desired direction. This has led recent research to behavioral analysis [21]. The behavioral analysis investigates social aspects like belief, attitude, intention, etc to explain individual behaviors with regard to misinformation [47]. Using reputation theory, [25] found that confirmation bias, source ratings, and argument quality, which are moderated by the veracity of information, had a significant positive effect on believability, and believability had a further positive effect on activities like commenting, and sharing. By exploiting the emotional content of tweets, [23] found that individuals are more likely to engage with negative information. This is supported by research that suggests that information reflecting negative emotions such as anger, fear, sadness, or doubt are spread more readily from person to person and more widely through social media than posts expressing joy [37]. Further research has shown that negativity boosts the likelihood of one's content being shared, particularly more for verified users on Twitter than ordinary users [41]. Additionally, an examination of the relationship between experiencing emotions and believing false news found that heightened emotionality is associated with an increased likelihood of belief in fake news and a decreased ability to differentiate between True and false news posts [31].
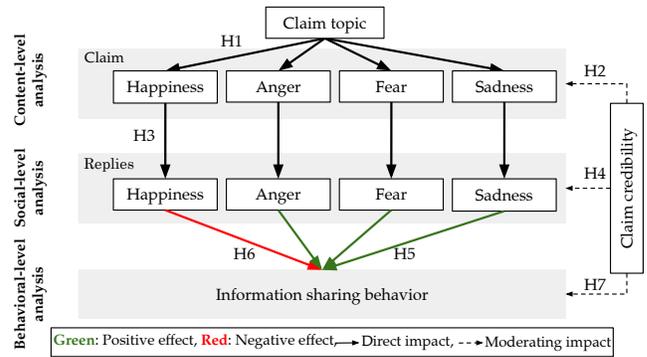

Figure 2: Proposed Research Model

Although the existing studies have revealed the significance of emotions in false claims and their influence on individuals' inclination to pass on information on social media, they fail to explain how these emotions fuel the greater spread of false claims compared to true ones. This paper contributes stronger empirical evidence compared to prior research in order to gain a comprehensive understanding of the role of emotions in the spread of false claims. Building on prior findings that suggest negative emotions are often associated with false claims, our research seeks to generalize these conclusions across various topics. Moreover, we analyze audience response posts to these claims to gauge their emotional reactions, and finally, we assess the influence of such emotions on their sharing behavior.

3 RESEARCH MODEL

Building upon previous studies presented in the literature review, to yield an interpretable analysis of information sharing behavior in social media, we set out a three-level research model underlying our study in Fig. 2. Details about the phases sketched in the research model have been explained in the following:

Content-level Analysis. At the first level of analysis, our goal is to understand how misinformation is expressed from an emotional perspective. For this purpose, we explore the correlation of emotions expressed in the claims with the claims' topic and credibility. To this end, the hypotheses developed are as follows:

- **H1.** The topic of a claim has a direct impact on the emotional language used in a claim.
- **H2.** Credibility of the claim has a moderating effect on the strength of emotions expressed in a claim.

Social-level Analysis. At the second level of analysis, we investigate how the emotional framing of topics in claims with different credibility influences the emotion of response posts by social media users to the claims. Indeed, we want to observe the reactions of readers by studying the emotions that they express in their responses. The underlying foundation of our idea is based on the assumption that the main emotion in the claims has a positive direct effect on the emotions of users expressed in their posts responding to these claims. The above arguments lead to the following hypotheses:
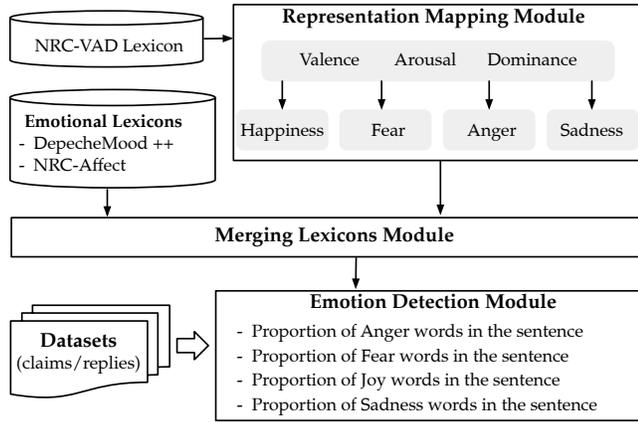
Figure 3: Overview of the framework for measuring emotions in claims and responses

- **H3.** The level of dominant emotion (happiness, anger, fear, and sadness) in a claim has a direct effect on the level of users' emotions expressed in their posts responding to the claim.
- **H4.** A claim's credibility has a moderating effect on emotions of response posts.

Behavioral-level Analysis. Finally, the third level of analysis investigates the influence of users' emotional responses related to false and true claims on sharing behavior in social media. The underlying assumption at this level of analysis is that, along with the topic, the emotion that was aroused in users towards a topic is an important factor in modeling activities such as retweeting and liking a claim. Thus, we explore the following hypotheses:

- **H5-6** Claims that trigger high levels of anger, fear, and sadness among users have a higher potential to be shared on social media as compared to claims that trigger a high level of happiness.
- **H7.** The credibility of the claims has a moderating effect on the sharing level of the questionable claims in social media.

4 MEASURING EMOTIONS METHODOLOGY

The computational framework of our methodology for measuring emotions is illustrated in Fig. 3. It includes three main modules, named Representation Mapping Module, Merging Lexicons Module, and Emotion Detection Module. The Representation Mapping and Merging Lexicons modules describe the approach that creates a unified emotion lexicon with high word coverage, enriched with the emotional intensity of words. The NRC Affect lexicon [34] serves as the basis for the joint representation. This lexicon was manually annotated with four basic emotions - *anger, fear, sadness, joy* - for both common English terms and terms that are prominent on social media platforms. The DepecheMood++ [2] and NRC-VAD [33] lexica are mapped into the representation used by NRC Affect, and the results are merged. The Emotion Detection Module extracts emotion vectors using the merged lexicon. The next sections provide details about the phases sketched in Fig. 3.

4.1 Representation Mapping Module

The presence of emotional words in a sentence provides a good premise for the interpretation of the overall emotion of the sentence [19]. Therefore, our first step towards emotion detection is discovering keywords and phrases that associate with emotion. Table 1 summarizes well-accepted resources used for acquiring basic knowledge about emotion-carrier words.

We investigated how similar to each other and different from each other these resources are. Table 2 shows the number of common words between each pair of lexicons, and also the number of words from each lexicon that have not been found in the remaining lexicons. The high number of unique words indicates that high coverage of emotional words can best be achieved by integrating lexicons. However, their integration is non-trivial, as the different resources are based on different representations. For example, the representation of the word *deadly* appears in different ways in the three lexicons (see Table 3). In DepecheMood++ and NRC-Affect *fear* has the highest score, and in the NRC-VAD lexicon, arousal has the highest score.

Exploring different lexicons revealed the need for a unified representation. To deal with different representations, we have developed the representation mapping module that maps the VAD representations of words in the NRC-VAD lexicon into ratings for four emotion labels (i.e. anger, fear, sadness, happiness), which is a joint representation covered by both the NRC-Affect and DepecheMood++ lexicons. According to Russell and Mehrabian [39] discrete affective states can be represented by the VAD dimensions. Table 4 presents the corresponding values of the VAD model for *happiness, anger, fear* and *sadness* emotions. Employing Table 4, we determined a four-dimensional emotional model of NRC-VAD lexicon words by computing the cosine similarity between the words and *happiness, anger, fear* and *sadness* emotions, as shown in (1),

$$\text{Emotional Score}(w, i) = \frac{\text{VAD}_w \times \text{VAD}_i}{\|\text{VAD}_w\|\|\text{VAD}_i\|} \tag{1}$$

$$i \in \{\text{Happiness, Anger, Fear, Sadness}\}$$

This equation assigns an emotional score to every word w in the NRC-VAD lexicon based on four emotional dimensions of *happiness, anger, sadness,* and *fear*. To avoid different label spaces that exist in different lexicons, before merging lexicons, values have been rescaled linearly such that the simple sum of them equals 1.

4.2 Merging Lexicons Module

To achieve maximum coverage of emotion carrying words, the Merging Lexicons Module integrates NRC-Affect, DepecheMood++ and NRC-VAD into our joint lexicon. NRC-Affect constitutes the basis to which words from DepecheMood++ are added if they are not available in NRC-Affect. To add a word from DepecheMood++ its 8-dimensional representation is shortened to consider only its happiness (=joy), fear, anger, and sadness scores. Finally, words that are in NRC-VAD, but not yet represented, are added based on the mapping from equation 1.

Table 1: Summary of the computational resources used in this paper.

Lexicon	Description
DepecheMood ++	DepecheMood ++ [2] is a high-precision and high-coverage emotion lexicon that has been automatically derived from crowd annotated news. It provides eight values between 0 and 1 ($e \in [0,1]^8$ and $\|e\| = 1$) for 62,224 entries on the following dimensions: *anger, anticipation, disgust, fear, joy, sadness, surprise,* and *trust.*
NRC-Affect	NRC Affect Intensity lexicon [34] is a list of 4,192 English words and their associations with four basic emotions (*anger, fear, sadness, joy*) with scores ranging from 0 to 1 ($e \in [0,1]^4$). It includes common English terms as well as terms that are prominent in social media platforms.
NRC-VAD	NRC Valence (positive–negative), Arousal (excited–calm), and Dominance (powerful–weak) lexicon [33] includes a list of more than 20,000 English words mapped to a 3D vector of VAD values, ranging from 0 (lowest) to 1 (highest) ($e \in [0,1]^3$).

Table 2: Number of unique words in each lexicon and number of words that overlap between each pair of lexicons.

Lexical resource	#Words	#Unique words	#Common words		
			DepecheMood ++	NRC-Affect	NRC-VAD
DepecheMood ++	62,224	54,456	-	1,790	7,759
NRC-Affect	4,192	0	1,790	-	4,192
NRC-VAD	20,006	9,740	7,759	4,192	-

Table 3: Representation of the word *deadly* in lexicons.

Lexical resource	Representations			
DepecheMood++	Anger	Fear	Sadness	Happiness
	0.17	0.36	0.21	0
NRC-Affect	Anger	Fear	Sadness	Happiness
	0.76	0.90	0.88	0
NRC-VAD	Valence	Arousal		Dominance
	0.14	0.85		0.55

Table 4: Values for *happiness, anger, fear* and *sadness* emotions in terms of VAD emotion dimensions [39].

Emotion	Valence	Arousal	Dominance
Happiness	0.76	0.48	0.35
Anger	-0.51	0.59	0.25
Fear	-0.64	0.60	-0.43
Sadness	-0.63	-0.27	-0.33

4.3 Emotion Detection Module

Emotion detection refers to the task of automatically assigning an emotion score to text from a set of predefined emotion labels [1]. Emotion lexicons are commonly used resources in unsupervised techniques to automatically and straightforwardly label text with emotional information, known as the keyword-based approach [3, 4, 29]. To discover emotions, we employ the unified emotion lexicon described in Section 4.2, which is based on DepecheMood++, NRC Affect, and NRC-VAD lexical resources. The "Emotion Detection Module" relies on the statistical features of emotional words to produce a vector of real-valued numbers as the result of the emotion assessment.

In the pre-processing phase, non-semantic words such as prepositions, conjunctions, and pronouns also known as stop words are removed from the input text. For out-of-vocabulary words, we used a difflib library in python to find the closest word to a given OOV word in the final joint lexicon. If no word with more than 90% of similarity is found, the word is determined as neutral with a default value [53]. We define proportions of happiness, anger, fear, sadness, or neutrality by extending our four-dimensional model to a five-dimensional. Now, given a sequence of words $s = (w_1...w_k)$ we can calculate the proportion for any of the five dimensions by

$$\text{Proportion (s, i)} = \frac{1}{k} \sum_{j=1}^{k} \pi_i(w_j) \qquad (2)$$

$$i \in \{\text{Happiness, Anger, Fear, Sadness, Neutral}\}$$

where π_i is the projection of the j-th component for emotion *i*. To better interpret the measured values in formula (2), we normalized emotional values according to the corresponding vector. Indeed, each value in the vector is located between 0 and 1, and the sum of the values for a sentence is equal to 1. Fig. 1 shows an example of emotion measurement using statistical features of emotional words.

5 DATA ANALYSIS AND RESULTS

In this section, we provide an assessment of our hypotheses on four recently released real-world data sets. Data analysis was performed using Structural Equation Modelling (SEM) with SmartPLS v3.3.9. The reason we used Partial Least Squares (PLS) was related to the research model complexity with moderating variables [18]. In addition, PLS accommodates non-normally distributed data. Our measurements were not normally distributed, both skewness and

Table 5: Data sets characteristics

Datasets	Domain	#Claims	#Replies
FA-KES [40]	Syrian War	752	-
FakeHealth [8]	Health	1,366	-
CoAID [7]	COVID-19	1,892	47,547
Co-inform [10]	Political	884	48,349

kurtosis were significant, which indicates that PLS is suited to this study [16]. Following the two-step analytical procedure suggested by [18], to estimate and analyze path coefficients, and test the proposed research hypotheses, we first assessed the measurement model, then, performed structural model analysis.

5.1 Data sets

The challenging problem of analyzing false claims is confined so far to one particular domain mostly with a focus on the field of politics. It was our objective to investigate the influence of topics on the emotional framing of the claims therefore we selected data to cover different topics. Secondly, data sets must contain the annotated claims as false and true along with users' responses to the claims. We evaluate our hypotheses across 4 recently released real-world data sets from four topics: politics, health, Syrian war, and COVID-19. To have a reliable analysis of results, in all data sets we selected a balanced number of false and true claims. Table 5 summarizes details of data sets used in this paper.

- **FA-KES** [40] data set comprises claims labeled as true or fake that cover the many facets of the Syrian war from the year 2013 to 2017. The annotation was supported by employing a semi-supervised fact-checking approach with the help of crowd-sourcing. To validate the accuracy of the labeling approach, a set of 50 claims was also manually assessed by domain experts.

- **FakeHealth** [8] data set is collected from 2005 to 2018 from Health News Review (www.healthnewsreview.org) website which assesses health news in aspects such as overclaiming, missing information, reliability of sources, and conflict of interests. The rating ranges from 0 to 5, and it is in proportion to the number of criteria satisfied by the claim. Following the strategy in [43], claims whose scores were lower than 3 were considered false.

- **CoAID** [7] (**C**ovid-19 he**A**lthcare m**I**sinformation **D**ata set) includes COVID-19 related claims, ranging from December 2019 through November 2020. Annotation has been made through fact-checking websites where the credibility of the claims is assessed by experts, which makes the quality of annotation high. Users' social engagement was also collected from Twitter which included the tweets discussing the claim in question and replies.

- **Co-inform** [10] data set contains political claims collected during the year 2020 from different fact-checking websites - such as PolitiFact, Factcheckni (www.factcheckni.org), FullFact (www.fullfact.org) and labeled as credible, when the fact-checkers detect it as true, or not-credible when the fact-checkers detect the claim is false. The collected claims in this data set check if the claim appeared on Twitter. This filter helped us to collect social context related to the claims from Twitter.

FA-KES and FakeHealth data sets contain only annotated claims as false and true, but, CoAID and Co-inform data sets in addition to the claims contain also responses, like, and retweets related to the claims.

5.2 Measurement Model Analysis

We performed measurement model validation, using SmartPLS v3.3.9, on data sets to assess the reliability and validity (convergent and discriminant validity) of constructs before testing hypotheses.

To this end, first, the *measurement reliability* is verified using composite reliability (CR) and Cronbach's alpha (CA) for each of the constructs [5, 6]. Both values exceeded the recommended threshold value of 0.7 [20] for all constructs, which can thus be regarded as reliable. Second, the *convergent validity* of constructs was assessed using the average variance extracted (AVE) [12]. As expected, all constructs showed satisfactory values above the suggested threshold value of 0.50. Thus, our study met the requirement for convergent validity. Third, to determine the *discriminant validity* of the measurement model, we assessed the Fornell-Larcker criterion of the constructs [12]. Since the square root of AVE for each construct exceeded the corresponding inter-construct correlations, thereby, the discriminant validity of the constructs for the constructs was supported. Overall, the results demonstrated the high reliability and validity of the measurement model.

5.3 Structural Model Analysis

This section proceeds to test the proposed research model. The PLS-SEM approach with a bootstrapping procedure was used to assess hypotheses. The key criteria for evaluating the structural model and the strengths of the relationships between the variables are to examine the t-values, and path coefficient (β values) [17]. The following sections provide the results of hypotheses testing on content-, social-, and behavioral-level.

5.3.1 Content-level Analysis. This section evaluates hypotheses H1 and H2 at the content-level of our research model, particularly, whether claims with different credibility elicit emotions with similar distributions on different topics. Our goal, in the first-level of analysis, is to understand if the basic word usage patterns differed from an emotional perspective in false and true claims and in different topics.

Fig. 4 provides a graphical comparison of emotions expressed in the FA-KES data set. Chart A shows the average of anger, fear, sadness, happiness, and neutrality in true and false claims. As one can see, in general, the claims convey negative emotions to their audience. Meanwhile, false claims show a higher intensity of "fear" compared to true claims, which makes sense since the topic of the claims is war. To have a better overview of emotions distribution in this data set, chart B depicts the emotional patterns in true and false claims. We call the distribution of scores over our four different emotions and neutrality the emotional pattern of a claim. The x-axis shows false and true claims and the y-axis shows the level of different emotions in claims. False and true claims demonstrate different emotional patterns. In true claims, you can see a similar distribution of emotion scores between anger, fear, and sadness. True claims tend to be more neutral than false claims, however, false claims tend to convey fear. Therefore, using the emotional

Table 6: Structural model results for hypothesis 1.

Hypothesised Relationship	Path Coefficient(t-value)
Claim Topic $\rightarrow$ Claim Anger	0.37 (5.16)***
Claim Topic $\rightarrow$ Claim Fear	0.25 (4.53)***
Claim Topic $\rightarrow$ Claim Sadness	0.17 (2.95)**
Claim Topic $\rightarrow$ Claim Happiness	0.01 (0.88)ns
Claim Topic $\rightarrow$ Claim Neutral	0.05 (1.22)ns

Note: * Significant at $p < 0.05$, ** at $p < 0.01$, and *** at $p < 0.001$.

pattern diagram it is predictable that, in the war-related claims, the level of fear expressed in claims is a signal that helps to assess their credibility.

Fig. 5 shows the correlation of emotions expressed in claims across the FakeHealth data set. Chart A shows the average score of different emotions in the claims. We can see that there is a balance between negative emotions in false and true claims, while the average of "happiness" expressed in false claims compared to true claims is higher. Chart B shows the scatter diagram that compares happiness with fear, sadness, anger, and neutrality in claims. Here it is more clear that happiness has higher intensity compared to other emotions in false claims. Detailed observations have shown that most of the false claims transmit intense "happiness" by promising an outstanding medicine or approach to cure diseases like cancer and diabetes, which is understandable since companies try to advertise and make a sale of their products.

Fig. 6 shows the average of different emotions expressed in false and true claims in the CoAID (chart A) and Co-inform (chart B) data sets. In true claims related to COVID-19, there is a balance among emotions, as well as a tendency for neutrality compared to false claims. While, in false claims, an intense inclination towards fear and then anger can be observed. In the Co-inform data set, containing US political claims, as we can see, regardless of the claims' credibility, negative emotions (i.e. anger, sadness, and fear) are dominant. Although in false claims the intensity of anger is higher than true claims.

Table 6 illustrates the path coefficient, t-values, and significance, for hypothesis H1. Hypothesis H1 addresses the structural relationships between the topic of a claim and the emotion expressed in a claim. To evaluate hypothesis H1, we combined claims from the four data sets and considered the topic of each data set as the topic of claims. As hypothesized, results reveal that the topic of the claim significantly influences the emotional language used in the claim. As is evident, the topic of the claims significantly affects the level of anger ($\beta = 0.37$; $p < 0.001$), fear ($\beta = 0.25$; $p < 0.001$), and sadness ($\beta = 0.17$; $p < 0.01$) expressed in the claims. Overall, this provides support for hypothesis H1. The relationship between the topic of the claim and the level of happiness expressed in the claims was not found significant ($p > 0.05$).

The path coefficient, t-values, and significance, for the hypothesis H2 testing results in each data set, are presented in Table 7. The results confirm that the credibility of the claims has a significant effect on the emotions expressed in the claims, offering support for H2. For example, in the FA-KES data set, containing claims related to the Syrian war, the credibility of the claims illustrates a positive direct effect on emotions, with the exception of fear ($\beta = -0.79$; $p < 0.001$). This means that the claims' credibility has a reverse significant effect on the level of fear in the claims, approving our observation in Fig. 4, that false claims, in this data set, have a higher intensity of "fear" compared to true ones.

Overall, the emotional language analysis of false and true claims on different topics provides evidence supporting our hypotheses that the intention of the claims' authors to incite emotion in their audience and the topic of the claims are inter-dependent and are influenced by the credibility of the claims.

5.3.2 Social-level Analysis. This section evaluates Hypotheses 3 and 4 at the social-level, to understand the influence of emotions expressed in the claims and also credibility of the claims on social media users' response posts.

Fig. 7 provides a comparison between the aroused emotions in users' responses to the claims on CoAID (chart A) and Co-inform (chart B) data sets by considering the claims' credibility. In corresponding replies to COVID-19 claims, the reflection of the emotions expressed in the claims (see Fig. 6 chart A) can be seen in users' responses. Emotion analysis of users' response posts to false claims related to COVID-19 and the global pandemic shows users' reactions, on average, have a higher intensity of fear and anger than responses to true claims. In the Co-inform data set, a similar distribution of emotion scores has been observed in users' responses and in the claims (see Fig. 6 chart B), and users' dominant emotional reaction in the responses is anger which is also the major emotion in the claims.

To have a closer look at the effect of emotions in the claims related to COVID-19 and their audiences' responses on social media, we partitioned the claims based on the main emotions in them (i.e., fear, anger, sadness, and happiness). Fig. 8 shows emotional pattern diagrams in claims with the highest fear, anger, sadness, and happiness scores and users' corresponding response posts on the right side in the CoAID data set. As is evident the emotional patterns in claims repeat themselves in users' response posts, and the dominant emotion in the claims directly affected users' emotions.

Tables 8 and 9 show the testing results of hypotheses 3 and 4 in CoAID and Co-inform data sets. As hypothesized, the results indicate the dominant emotion in a claim has a significant effect on users' emotions expressed in their response posts, thus verifying H3. Likewise, as expected, the credibility of the claims significantly affects emotions of response posts, offering support for H4.

5.3.3 Behavioral-level Analysis. This section evaluates Hypotheses 5, 6, and 7 at the behavioral-level, particularly, the impact of emotions expressed in false and true claims and corresponding users' response posts on their *sharing* behavior in social media. Fig. 9 shows the average of retweets, likes, and replies in false and true claims in the CoAID and Co-inform data sets. In the CoAID data set, we can see that the average of retweets and likes in false claims is significantly higher than true ones, while, the average number of users' response posts to true claims is overtaking false claims. In the Co-inform data set, the tendency for behaviors like retweeting and responding is slightly higher for false claims than for true claims.

To determine the emotions that caused the highest levels of activity by users, we counted retweets, likes, and replies depending on dominating emotion. To this aim, the emotion with the highest
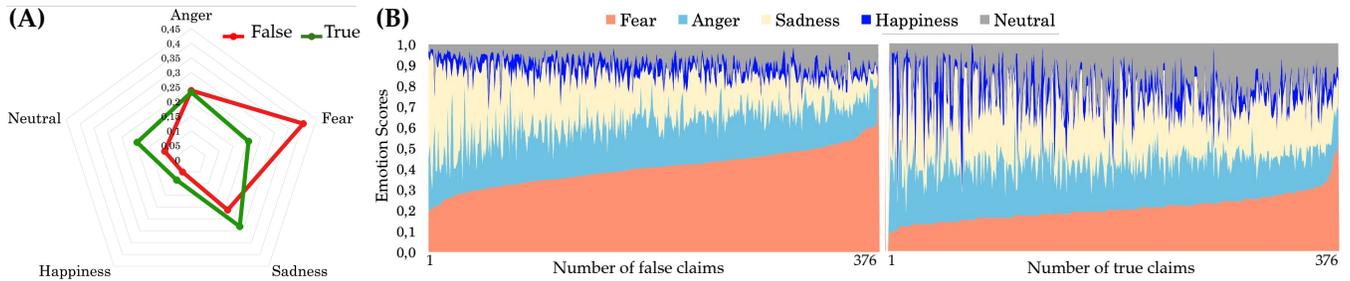
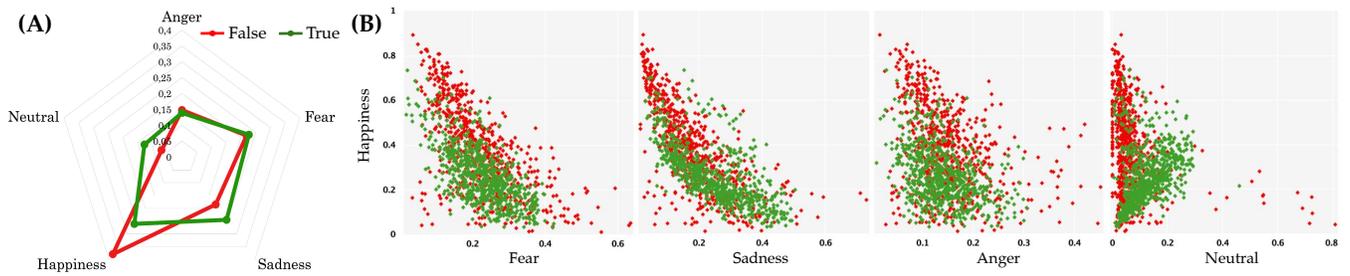**Figure 4: The averaged emotion scores (A) and emotional pattern diagrams (B) in FA-KES data set.**

Figure 5: The averaged emotion scores (A) and Scatter plots of happiness intensity vs fear, sadness, anger, and neutral (B) in FakeHealth data set.

Table 7: Structural model results of hypothesis 2.

Hypothesised Relationship	Path Coefficient(*t*-value)			
	FA-KES	FakeHealth	CoAID	Co-inform
Claim Credibility → Claim Anger	0.20 (6.82)***	0.07 (2.70)**	-0.32 (8.37)***	-0.21 (4.29)***
Claim Credibility → Claim Fear	-0.79(23.65)***	0.05 (1.63)ns	-0.30 (6.53)***	0.16 (3.95)***
Claim Credibility → Claim Sadness	0.54 (16.10)***	0.27 (10.24)***	0.14 (3.82)***	-0.08 (1.16)ns
Claim Credibility → Claim Happiness	0.06 (1.77)ns	-0.34 (15.03)***	0.28 (3.62)***	0.02 (0.43)ns
Claim Credibility → Claim Neutral	0.21 (4.12)***	0.35 (8.64)***	0.15 (3.57)***	0.19 (2.61)**

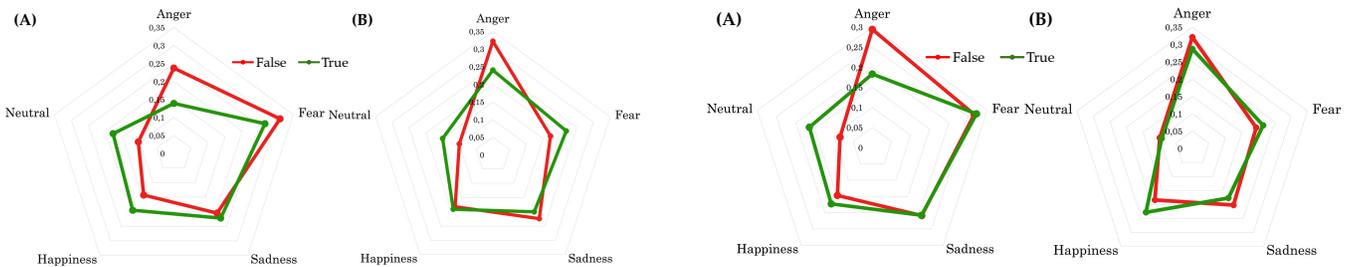Note: * Significant at $p < 0.05$, ** at $p < 0.01$, and *** at $p < 0.001$.

Figure 6: The average of emotion scores across true and false claims in CoAID (A) and Co-inform (B) data sets.

Figure 7: The average of emotion scores in replies to true and false claims in CoAID (A) and Co-inform (B) data sets.

score in a claim has been defined as its main emotion, and claims have been partitioned according to their main emotion. Tables 10 and 11 show the levels of activity in CoAID and Co-inform data sets depending on dominant emotions and credibility of claims. In

the CoAID data set, the highest levels of activity were observed for claims expressing anger. Furthermore, negative emotions in false claims, especially anger and sadness, spurred higher numbers of retweets and likes than in true claims. In the claims expressing
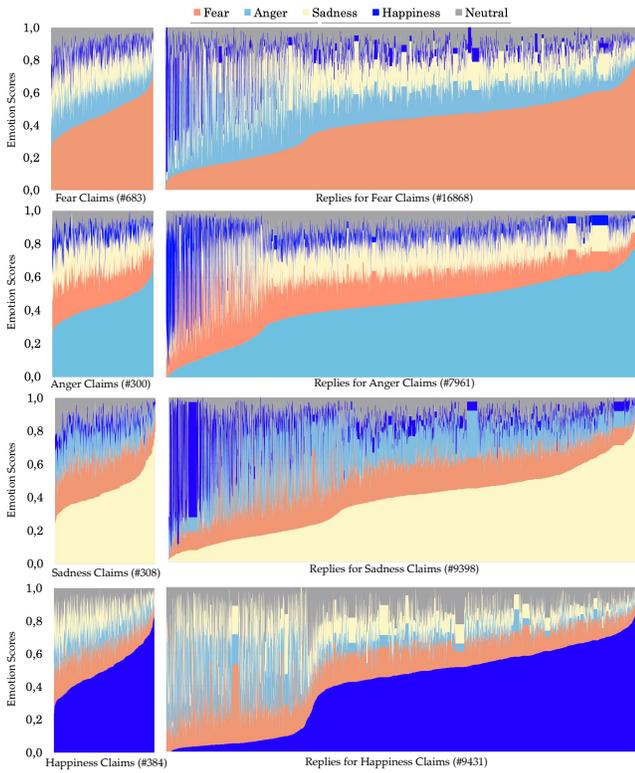
Figure 8: Emotional pattern diagram in the claims (left) and their related replies (right) in CoAID data set.

Table 8: Structural model results of hypothesis 3.

Hypothesised Relationship	Path Coefficient(t-value)	
	CoAID	Co-inform
Claim Anger → Replies Anger	0.51 (11.19)***	0.44 (11.47)***
Claim Fear → Replies Fear	0.57 (15.72)***	0.43 (9.40)***
Claim Sadness → Replies Sadness	0.57 (12.71)***	0.39 (8.30)***
Claim Happiness → Replies Happiness	0.58 (12.98)***	0.38 (8.03)***
Claim Neutral → Replies Neutral	0.32 (6.83)***	0.35 (5.38)***

Note: * Significant at $p < 0.05$, ** at $p < 0.01$, and *** at $p < 0.001$.

happiness, true claims have received higher retweets and likes from users than false claims. Similarly, in the Co-inform data set, in the claims carrying negative emotions, false claims caused higher retweets, like, and responses in users, and in the claims expressing happy emotions, true claims triggered a higher level of interaction between users.

The path coefficient, t-values, and significance, for hypotheses 5, 6, and 7 testing results in CoAID and Co-inform data sets are tabulated in Table 12. In the CoAID data set, the results indicate that the level of anger ($\beta = 0.44$, $p < 0.001$), fear ($\beta = 0.22$, $p < 0.05$), and sadness ($\beta = 0.28$, $p < 0.01$) in users' replies toward the claims have a significant effect on the sharing of claims. Likewise, in the Co-inform data set, negative emotions (anger, fear, and sadness)

Table 9: Structural model results of hypothesis 4.

Hypothesised Relationship	Path Coefficient(t-value)	
	CoAID	Co-inform
Claim Credibility → Replies Anger	-0.13 (3.99)***	-0.14 (3.75)***
Claim Credibility → Replies Fear	-0.11 (3.65)***	0.03 (2.09)*
Claim Credibility → Replies Sadness	0.08 (2.43)*	-0.10 (2.05)*
Claim Credibility → Replies Happiness	0.10 (2.92)**	0.01 (3.10)**
Claim Credibility → Replies Neutral	0.21 (5.14)***	0.08 (3.90)***

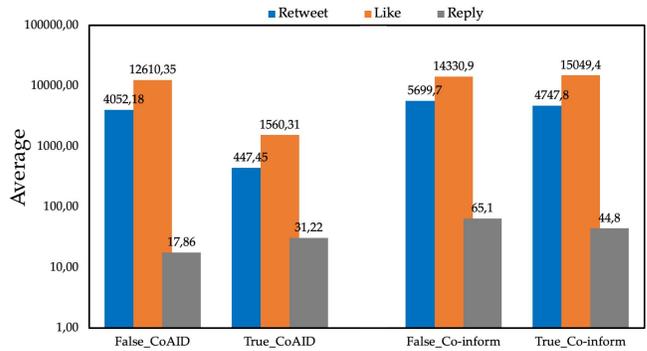Note: * Significant at $p < 0.05$, ** at $p < 0.01$, and *** at $p < 0.001$.

Figure 9: The average of retweets, likes and replies across true and false claims in the CoAID and Co-inform data sets.

Table 10: Average of retweets, likes and replies separated by credibility and emotions in CoAID data set.

Emotion	Credibility (#Claims)	Retweet	Like	Reply
Anger	False (265)	**11987.5**	**33383.0**	19.2
	True (35)	2436.2	7943.9	**82.2**
Fear	False (398)	**1095.6**	**5389.6**	15.8
	True (285)	728.0	2551.2	**37.2**
Sadness	False (89)	**2157.1**	**9229.6**	4.1
	True (219)	190.8	595.2	**25.0**
Happiness	False (156)	177.7	734.2	9.0
	True (228)	**365.7**	**1456.2**	**35.2**
Neutral	False (38)	24.1	**49.3**	5.4
	True (179)	**30.1**	47.8	**15.7**

in users' replies have a significant effect on the sharing of claims. Thus, hypothesis 5 is supported. Meanwhile, the level of happiness in users' replies to claims, does not have a significant effect on the sharing behavior; thus, hypothesis H6 was not supported. In addition, neutrality in users' replies does not affect sharing behavior ($p > 0.05$). Moreover, our analysis shows that the credibility of the claims has a negative direct effect on sharing of questionable claims ($\beta = -0.27$; $p < 0.001$) and ($\beta = -0.05$; $p < 0.01$) in the CoAID and Co-inform data set, respectively. Therefore, hypothesis 7 is supported, meaning that claims' credibility significantly influences sharing them. The PLS results show that the variance explained for sharing behavior was approximately 52% and 41%, in CoAID

Table 11: Average of retweets, likes and replies separated by credibility and emotions in Co-inform data set.

Emotion	Credibility (#Claims)	Retweet	Like	Reply
Anger	False (231)	**6236.6**	14854.7	**66.5**
	True (152)	5199.6	**17424.5**	54.3
Fear	False (39)	**9607.7**	**22772.4**	**65.4**
	True (77)	6355.1	12699.9	47.3
Sadness	False (67)	**5286.4**	**14519.1**	**61.2**
	True (60)	2956.4	8992.7	33.7
Happiness	False (84)	2829.9	8785.9	**61.0**
	True (97)	**5951.6**	**23428.0**	40.9
Neutral	False (24)	**5379.0**	**14454.0**	**69.5**
	True (53)	941.6	3173.5	30.1

Table 12: Structural model results of hypotheses 5, 6, and 7.

Hypothesis	Hypothesised Relationship	Path Coefficient(t-value)	
		CoAID	Co-inform
H5-H6	Replies Anger → Retweet	0.44 (5.80)***	0.31 (5.34)***
	Replies Fear → Retweet	0.22 (1.99)*	0.38 (7.72)***
	Replies Sadness → Retweet	0.28 (2.99)**	0.12 (2.71)**
	Replies Happiness → Retweet	0.08 (0.66)ns	0.11 (2.12)*
	Replies Neutral → Retweet	0.04 (0.34)ns	0.07 (0.13)ns
H7	Claim Credibility → Retweet	-0.27 (10.91)***	-0.05 (2.60)**

Note: * Significant at $p < 0.05$, ** at $p < 0.01$, and *** at $p < 0.001$.

and Co-inform data sets, respectively. Overall, the results provide evidence supporting our hypothesised model.

6 CONCLUSION

Studies that investigate the behavior of social media users with regard to misinformation indicate that emotions in false claims play a crucial role in spurring people's willingness to share information on social media. However, these studies do not explain how emotions enhance the potency of false claims, causing them to spread more widely than true claims. In this paper, we have presented an in-depth analysis of the effects of emotional framing in false and true claims from different topics on the users' emotions expressed in replies and distorting their sharing behavior. The emotional language analysis of the claims showed that the distribution of emotions can be influenced by both the topic and the credibility of the claims. The deeper analysis of the results showed the reflection of the emotions expressed in the claims in users' emotional responses on social media. Furthermore, our diffusion analysis suggests that not only the topic of the claim itself but also the emotion related to the topic is an effective subtle driving factor in users' retweeting decision. These findings suggest that social media users need to be aware of the potential influence of emotional appeals and the credibility of claims when sharing information online.

Additionally, we found that the ability of emotions to influence sharing behaviors on social media depends on the specific emotion that was aroused. False claims motivate the audience to retweet the claims via negative emotional appeals such as anger, fear, and sadness. While, in claims that arouse happiness in the audience,

true claims caused more interaction among users than false claims. In the future, we aim to explore the relationship between emotional framing and emotional contagion in online discussions. We will investigate the emotional contagion changes in online discussions over time, and the emotional contagion effect on spreading truth and falsity.

REFERENCES

[1] Nourah Alswaidan and Mohamed El Bachir Menai. 2020. A survey of state-of-the-art approaches for emotion recognition in text. *Knowledge and Information Systems* (2020), 1–51.
[2] Oscar Araque, Lorenzo Gatti, Jacopo Staiano, and Marco Guerini. 2019. Depechemood++: a bilingual emotion lexicon built through simple yet powerful techniques. *IEEE transactions on affective computing* (2019).
[3] Christos Baziotis, Nikos Athanasiou, Alexandra Chronopoulou, Athanasia Kolovou, Georgios Paraskevopoulos, Nikolaos Ellinas, Shrikanth Narayanan, and Alexandros Potamianos. 2018. Ntua-slp at semeval-2018 task 1: Predicting affective content in tweets with deep attentive rnns and transfer learning. *arXiv preprint arXiv:1804.06658* (2018).
[4] Erik Cambria, Dipankar Das, Sivaji Bandyopadhyay, Antonio Feraco, et al. 2017. A practical guide to sentiment analysis. (2017).
[5] Wynne W Chin et al. 1998. The partial least squares approach to structural equation modeling. *Modern methods for business research* 295, 2 (1998), 295–336.
[6] Lee Anna Clark and David Watson. 2016. Constructing validity: Basic issues in objective scale development. (2016).
[7] Limeng Cui and Dongwon Lee. 2020. CoAID: COVID-19 Healthcare Misinformation Dataset. *arXiv preprint arXiv:2006.00885* (2020).
[8] Enyan Dai, Yiwei Sun, and Suhang Wang. 2020. Ginger Cannot Cure Cancer: Battling Fake Health News with a Comprehensive Data Repository. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 853–862.
[9] Michela Del Vicario, Alessandro Bessi, Fabiana Zollo, Fabio Petroni, Antonio Scala, Guido Caldarelli, H Eugene Stanley, and Walter Quattrociocchi. 2016. The spreading of misinformation online. *Proceedings of the National Academy of Sciences* 113, 3 (2016), 554–559.
[10] Ronald Denaux, Martino Mensio, Jose Manuel Gomez-Perez, and Harith Alani. 2021. Weaving a Semantic Web of Credibility Reviews for Explainable Misinformation Detection. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization.
[11] Carolyn S Fish and Amy L Griffin. 2019. Emotional responses to climate change map framing using facial emotion recognition technology. *Abstracts of the ICA* 1 (2019).
[12] Claes Fornell and David F Larcker. 1981. Evaluating structural equation models with unobservable variables and measurement error. *Journal of marketing research* 18, 1 (1981), 39–50.
[13] Maria Glenski, Tim Weninger, and Svitlana Volkova. 2018. Identifying and understanding user reactions to deceptive and trusted social news sources. *arXiv preprint arXiv:1805.12032* (2018).
[14] Jennifer Golbeck, Matthew Mauriello, Brooke Auxier, Keval H Bhanushali, Christopher Bonk, Mohamed Amine Bouzaghrane, Cody Buntain, Riya Chanduka, Paul Cheakalos, Jennine B Everett, et al. 2018. Fake news vs satire: A dataset and analysis. In *Proceedings of the 10th ACM Conference on Web Science*. 17–21.
[15] Chuan Guo, Juan Cao, Xueyao Zhang, Kai Shu, and Miao Yu. 2019. Exploiting emotions for fake news detection on social media. *arXiv preprint arXiv:1903.01728* (2019).
[16] Joseph F Hair, G Tomas M Hult, Christian M Ringle, Marko Sarstedt, and Kai Oliver Thiele. 2017. Mirror, mirror on the wall: a comparative evaluation of composite-based structural equation modeling methods. *Journal of the academy of marketing science* 45, 5 (2017), 616–632.
[17] JF Hair Jr, GTM Hult, Christian M Ringle, and M Sarstedt. 2017. A primer on partial least squares equation modeling (PLS-SEM).
[18] Joseph F Hair Jr and Marko Sarstedt. 2019. Factors versus composites: Guidelines for choosing the right structural equation modeling method. *Project Management Journal* 50, 6 (2019), 619–624.
[19] Akram Sadat Hosseini. 2017. Sentence-level emotion mining based on combination of adaptive Meta-level features and sentence syntactic features. *Engineering Applications of Artificial Intelligence* 65 (2017), 361–374.
[20] John Hulland. 1999. Use of partial least squares (PLS) in strategic management research: A review of four recent studies. *Strategic management journal* 20, 2 (1999), 195–204.
[21] Ming Jia, Hongfei Ruan, and Zhe Zhang. 2017. How rumors fly. *Journal of Business Research* 72 (2017), 33–45.
[22] Zhiwei Jin, Juan Cao, Yongdong Zhang, and Jiebo Luo. 2016. News verification by exploiting conflicting social viewpoints in microblogs. In *Thirtieth AAAI*

conference on artificial intelligence.

[23] Jari Kätsyri, Teemu Kinnunen, Kenta Kusumoto, Pirkko Oittinen, and Niklas Ravaja. 2016. Negativity bias in media multitasking: The effects of negative social media messages on attention to television news broadcasts. *PloS one* 11, 5 (2016), e0153712.

[24] Irina Khaldarova and Mervi Pantti. 2016. Fake news: The narrative battle over the Ukrainian conflict. *Journalism practice* 10, 7 (2016), 891–901.

[25] Antino Kim and Alan Dennis. 2018. Says who?: How news presentation format influences perceived believability and the engagement level of social media users. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.

[26] David Klein and Joshua Wueller. 2017. Fake news: A legal perspective. *Journal of Internet Law (Apr. 2017)* (2017).

[27] Rinaldo Kühne et al. 2014. Political news, emotions, and opinion formation: Toward a model of emotional framing effects. In *Annual Conference of the International Communication Association (ICA), Phoenix, AZ*.

[28] David F Larcker and Anastasia A Zakolyukina. 2012. Detecting deceptive discussions in conference calls. *Journal of Accounting Research* 50, 2 (2012), 495–540.

[29] Yukun Ma, Haiyun Peng, and Erik Cambria. 2018. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.

[30] A Mackey, Susan Gauch, and Kevin Labille. 2021. Detecting fake news through emotion analysis. In *Proceedings of the 13th International Conference on Information, Process, and Knowledge Management*. 65–71.

[31] Cameron Martel, Gordon Pennycook, and David G Rand. 2020. Reliance on emotion promotes belief in fake news. *Cognitive research: principles and implications* 5 (2020), 1–20.

[32] Rada Mihalcea and Carlo Strapparava. 2009. The lie detector: Explorations in the automatic recognition of deceptive language. In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*. 309–312.

[33] Saif Mohammad. 2018. Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 174–184.

[34] Saif Mohammad and Svetlana Kiritchenko. 2018. Understanding emotions: A dataset of tweets to study interactions between affect categories. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.

[35] Matthew L Newman, James W Pennebaker, Diane S Berry, and Jane M Richards. 2003. Lying words: Predicting deception from linguistic styles. *Personality and social psychology bulletin* 29, 5 (2003), 665–675.

[36] Babajide Osatuyi and Jerald Hughes. 2018. A tale of two internet news platforms-real vs. fake: An elaboration likelihood model perspective. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.

[37] Max Pellert, Simon Schweighofer, and David Garcia. 2020. The individual dynamics of affective expression on social media. *EPJ Data Science* 9, 1 (2020), 1.

[38] Kirk Roberts, Michael A Roach, Joseph Johnson, Josh Guthrie, and Sanda M Harabagiu. 2012. EmpaTweet: Annotating and Detecting Emotions on Twitter.. In *Lrec*, Vol. 12. Citeseer, 3806–3813.

[39] James A Russell and Albert Mehrabian. 1977. Evidence for a three-factor theory of emotions. *Journal of research in Personality* 11, 3 (1977), 273–294.

[40] Fatima K Abu Salem, Roaa Al Feel, Shady Elbassuoni, Mohamad Jaber, and May Farah. 2019. Fa-kes: A fake news dataset around the syrian war. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 13. 573–582.

[41] Jonas Schöne, David Garcia, Brian Parkinson, and Amit Goldenberg. 2022. Negative Expressions are Shared more on Twitter for Public Figures than for Ordinary Users. (2022).

[42] Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. 2019. Combating fake news: A survey on identification and mitigation techniques. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 3 (2019), 1–42.

[43] Kai Shu, Deepak Mahudeswaran, Suhang Wang, Dongwon Lee, and Huan Liu. 2020. FakeNewsNet: A Data Repository with News Content, Social Context, and Spatiotemporal Information for Studying Fake News on Social Media. *Big Data* 8, 3 (2020), 171–188.

[44] Kai Shu, Deepak Mahudeswaran, Suhang Wang, and Huan Liu. 2020. Hierarchical propagation networks for fake news detection: Investigation and exploitation. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 14. 626–637.

[45] Kai Shu, Suhang Wang, and Huan Liu. 2019. Beyond news contents: The role of social context for fake news detection. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*. ACM, 312–320.

[46] Qi Su, Mingyu Wan, Xiaoqian Liu, and Chu-Ren Huang. 2020. Motivations, Methods and Metrics of Misinformation Detection: An NLP Perspective. *Natural Language Processing Research* (2020).

[47] Sandeep Suntwal, Susan Brown, and Mark Patton. 2020. How does Information Spread? An Exploratory Study of True and Fake News. In *Proceedings of the 53rd Hawaii International Conference on System Sciences*.

[48] Sander Van Der Linden. 2017. The nature of viral altruism and how to make it stick. *Nature Human Behaviour* 1, 3 (2017), 1–4.

[49] Svitlana Volkova, Kyle Shaffer, Jin Yea Jang, and Nathan Hodas. 2017. Separating facts from fiction: Linguistic models to classify suspicious and trusted news posts on twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 647–653.

[50] Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359, 6380 (2018), 1146–1151.

[51] Xiaobao Wang, Di Jin, Mengquan Liu, Dongxiao He, Katarzyna Musial, and Jianwu Dang. 2019. Emotional contagion-based social sentiment mining in social networks by introducing network communities. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1763–1772.

[52] Yangqian Wang, Hao Han, Ye Ding, Xuan Wang, and Qing Liao. 2019. Learning Contextual Features with Multi-head Self-attention for Fake News Detection. In *International Conference on Cognitive Computing*. Springer, 132–142.

[53] Wojciech Witon, Pierre Colombo, Ashutosh Modi, and Mubbasir Kapadia. 2018. Disney at iest 2018: Predicting emotions using an ensemble. In *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. 248–253.

[54] Lianwei Wu and Yuan Rao. 2020. Adaptive interaction fusion networks for fake news detection. *arXiv preprint arXiv:2004.10009* (2020).

[55] Chaowei Zhang, Ashish Gupta, Christian Kauten, Amit V Deokar, and Xiao Qin. 2019. Detecting fake news for reducing misinformation risks using analytics approaches. *European Journal of Operational Research* 279, 3 (2019), 1036–1052.

[56] Xueyao Zhang, Juan Cao, Xirong Li, Qiang Sheng, Lei Zhong, and Kai Shu. 2021. Mining dual emotion for fake news detection. In *Proceedings of the web conference 2021*. 3465–3476.