Sensemaking of Socially-Mediated Crisis Information

Vrushali Koli, Jun Yuan, and Aritra Dasgupta
Department of Data Science, New Jersey Institute of Technology, USA
{vk435, jy448, aritra.dasgupta} @njit.edu

Abstract

In times of crisis, the human mind is often a voracious information forager. It might not be immediately apparent what one wants or needs, and people frequently look for answers to their most pressing questions and worst fears. In that context, the pandemic has demonstrated that social media sources, like erstwhile Twitter, are a rich medium for data-driven communication between experts and the public. However, as lay users, we must find needles in a haystack to distinguish credible and actionable information signals from the noise. In this work, we leverage the literature on crisis communication to propose an AI-driven sensemaking model that bridges the gap between what people seek and what they need during a crisis. Our model learns to contrast social media messages concerning expert guidance with subjective opinion and enables semantic interpretation of message characteristics based on the communicative intent of the message author. We provide examples from our tweet collection and present a hypothetical social media usage scenario to demonstrate the efficacy of our proposed model.

1 Introduction

During the early months of a crisis, people are not equipped with relevant knowledge about a crisis, such as what has occurred, what steps to take next, etc., and information can keep evolving rapidly. Public officials and crisis responders have often used social media to communicate crisis information (Graham et al., 2015). As witnessed during the pandemic, social media platforms not only shaped people's behavior and opinions but also served as a ground for communicating scientific information about public health.

It is widely acknowledged that conflicting information and claims can confuse the public, leading to counterproductive preventive actions, as seen during the COVID-19 pandemic (Rossmann et al., 2018). To the best of our knowledge there has not been much research focusing on designing or investigating social media sites when people might not know what they need (Jang and Baek, 2019) to navigate an unknown crisis. Some recent work investigates cognitive factors to identify the relationship between crisis type, organization reputation, and sentiments (Eaddy and Jin, 2018; Liu et al., 2016). The work in this domain explores distinct factors like forgiveness, empathy, anxiety etc, (Kim and Yang, 2009). Nonetheless, they do not focus on how people perceive crisis information and guide their decisions based on sensemaking (Stieglitz et al., 2017) and contextualization. Research in crisis communication suggests that effective communication requires an understanding of how different people perceive the messages, and what the fundamental drivers are for their information-seeking needs.

In our work, we take inspiration from the crisis communication literature for analyzing the different information facets that are needed by lay social media users to make sense of an unfolding, uncertain situation. Our work offers insight into the generalizability of crisis information-seeking characteristics. We contribute a crisis-related intent classification model that is eventually integrated into a human-AI interface to help social media users triage and group relevant information without being exposed to unnecessary noise and negativity. We demonstrate the efficacy of the model by describing a hypothetical sensemaking workflow of a social media user that leverages our proposed model and interfaces.

2 Related Work

We discuss the related work regarding two overlapping threads of research: i) socially mediated crisis communication and ii) AI-driven sensemaking using social media interfaces.

Proceedings of the Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing (3rd HCI+NLP Workshop), pages 74–81
June 21, 2024

2.1 Crisis Communication and Social Media

Studies have shown that the initial crisis stage is a critical period where informing people about protective behaviors supports a resilient crisis response (Islam et al., 2023; Bukar et al., 2022b,a). The use of social media bridged the gap between public officials and general public by providing a real-time communication platform.

The role of social media users has become more visible and active leading to collaborative crises response with public officials (Reuter and Kaufhold, 2018). As such, social media users rely on various sources during a crisis (Islam et al., 2023; Meadows et al., 2019). Current research on social media and crisis communication focuses on identifying influential personnel who can allow for efficient information dissemination.

Crisis communication aims to provide public with credible sources of information during the unfolding of a crisis (Lin et al., 2016). The information from trustworthy sources can help to curb the propagation of rumors (Aguirre and Tierney, 2001). Most messages on social media can be categorized as threat/risk messages or perceived severity (Myneni et al., 2023; Islam et al., 2023). Messages that include both emotional appeal and message source can impact how people perceive severity (Vaala et al., 2022). High engagement messages that are emotionally charged can affect the general public's response (Naseem et al., 2021). These messages help understand the impact of a crisis. However, the interplay among different information dimensions, like engagement, awareness, and cognitive load, remains unclear and is an active area of research (Stieglitz et al., 2017) that we contribute to.

2.2 AI-Driven Sensemaking of Socially Mediated Information

AI models are often used for distinguishing between facts and opinions on social media. The current work in fact-checking focuses on the automated classification of social media content using supervised learning algorithms. While these research studies present a fundamental approach towards identifying and solving the problem of check-worthy claim identification (Miranda et al., 2019; Hassan et al., 2017), the focus is only on binary classification tasks of accepting or refuting the claims (Hanselowski et al., 2019; Nakov et al., 2021). From the information consumers' perspective, fact-checking tools are out of user control as external sources provide them. Moreover, during a pandemic, users are generally navigating terra incognita as there is no establishment of ground truth that can be automatically detected. Users often want to be self-reliant and not completely rely on third-party fact-checking sites (Myneni et al., 2023). Currently, social media platforms allow end-users to curate their information feed by allowing users to filter what types of content they are exposed to. Studies distinguish between actions users can take to moderate content based on source (specific users) and types of content (Jhaver et al., 2023). Reducing types of content that are not exposed to the user can aid their information-seeking process by reducing the search space (Gillespie, 2022). Lack of transparency in algorithmic details of provided methods and unclear definitions of contextual terms can lead to further confusion. To address these challenges, we propose to give social media users the agency and control their feed while also carefully considering the role of AI-driven automation in triaging information that users might need but not necessarily be aware of owing to the uncertain information landscape.

3 Methodology

In this section, we describe our methodology for data collection and qualitative labeling of tweets that preceded the conceptualization of communicative intent. Please find details about the methodology in the supplemental material: `https://tinyurl.com/mrymxwed`.

Data Collection: We used the erstwhile Twitter User Timeline API (Hossain et al., 2018) for collecting tweets from March 2020 to September 2020. In particular, we wanted to collect tweets in two batches: one, focused on identifying regular social media users who could also be considered subject matter experts, and two, tweets from lay social media users. With the help of a published list from Elemental, a health and wellness publication, we collected the relevant COVID-19 tweets from a list of 50 health and science experts (Editors, 2020) who regularly updated information about the COVID-19 pandemic. With the aid of a researcher in the medical sciences domain, we verified that these people could be considered credible voices about the pandemic while acknowledging that there could be differences of opinion among experts. The second category of tweets in our collection is general

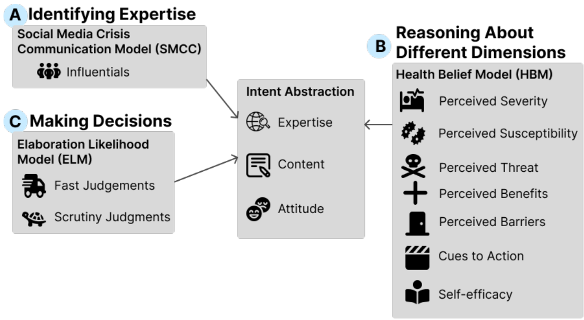

Figure 1: Illustrating the derivation of our Intent Model from the existing Crisis Communication Models (SMCC, ELM, and HBM).

user tweets. These come from users except these 50 experts over the same period. For both categories, expert and general users' tweets, we extract tweets using COVID-19-relevant keywords some of which include "coronavirus", "sars-cov-2", and "covid-19" and several others. We needed to ensure our general user tweet collection did not contain any tweets from users who could be considered experts. We ensured that general public tweets excluded tweets from users whose profiles included the keywords "epidemiologist", "virologist", "clinician", etc.

Topic Modeling. We anchored our analysis to understanding how experts and general users could well be discussing different dimensions of the pandemic. We leveraged publicly available deep-learning models with good performance on tweets based on BERT (Devlin et al., 2018) architecture for consistency across each information dimension. We use a transformer-based algorithm for topic modeling called BERTopic (Grootendorst, 2020). We trained the model on a subset of 5,000 expert tweets using c-TF-IDF to reduce outliers. Then, two graduate students labeled approximately 900 raw topic clusters to 25 topic categories using references from prior work on topic modeling covid-19 tweets (Oliveira et al., 2022; Vijayan, 2021; Lyu et al., 2021; Boon-Itt et al., 2020; Abd-Alrazaq et al., 2020). These topics include governmental affairs, vaccine development, scientific information, healthcare, mitigation, symptoms, etc.

Subjectivity and Sentiment Analysis. Subjectivity prediction can help consumers evaluate the text for more effective and efficient scientific communication. We, therefore collaborated with industrial researchers for using professional labeling services to tag 10,000 tweets from our corpus. The available labels were "objective, slightly objective, uncertain, slightly subjective, subjective, Irrelevant". We split the resulting data set into 3,232 for training and the remaining 808 samples for test tweets and trained a DistilBERT (Sanh et al., 2019) model for this task. Tweets labelled "uncertain" or "Irrelevant" were removed. We fine-tuned a pre-trained model to classify two labels. The total training time for the model is approximately one hour with the GPU-enabled Google Colab in the free tier. We expected expert messages to be more objective and general users to be more subjective. However, we found that experts also exhibited subjectivity in their tweets like general users across most topics.

For sentiment analysis, we chose the model bertweet-base-sentiment-analysis (Pérez et al., 2021) provided by HuggingFace. The model classification results in each tweet with three probabilities corresponding to positive, negative, and neutral. We use the label with the highest probability as the final label for the tweet.

4 Sensemaking via Communicative Intent

Communicative intent, or simply intent, refers to the aim or purpose of a tweet. Intent analysis can help information consumers determine whether the tweet is relevant to what they are seeking as it can provide contextual information about a particular topic. The message intent can be considered during the reasoning process (Monti et al., 2022) which can aid users in navigating the information space. Unprecedented emergencies, like the pandemic, require the public to adapt to time, domain, and context-specific information in understanding the communication dynamics on social media. Current crisis communication models are insufficient in guiding people when exposed to exponentially more information due to increased social media use. The intent classification model is shown in Figure 1. The related models to our work are the Health Belief Model (HBM), the Social Mediated Crisis Communication Model (SMCC), and the Elaboration Likelihood Model (ELM).

HBM states that an individual's personal beliefs affect their health-related behaviors (Washburn, 2020). This is a valuable framework to characterize people's discussions on social media based on perceived severity, perceived threat, perceived susceptibility, etc. However, this model does not give us a way to characterize why people perceive certain information in a particular way or what constitutes a threat. In the Intent Model, source,

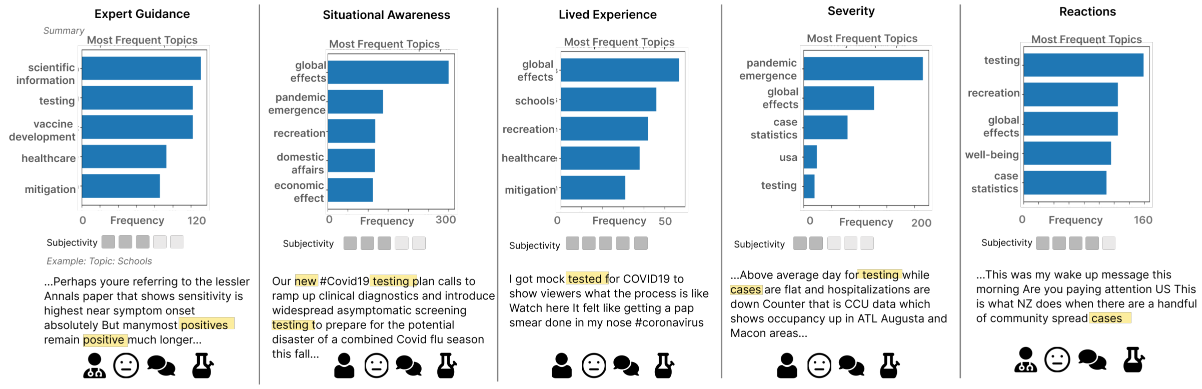

Figure 2: A snapshot of each intent category (1) An overview of each intent category showing the five most frequent topics and average subjectivity (2) An example tweet profiled by sentiment, subjectivity, and topic from both experts and general users.

content, and attitude are ways to characterize who, what, and how information is perceived under a particular category (Figure 1.2).

The SMCC model offers a solution by characterizing the source of information. The SMCC model emphasizes identifying individuals that can aid in information dissemination efforts, classifying these individuals as *influentials* (Liu et al., 2020). In social media sites, these individuals have high engagement by having a large following. However, these individuals may not be crisis experts, and other social media users may not want to see information from this individual. In the Intent Model this is characterized by source of information where we make a distinction between *crisis experts* and *general users* (Figure 1.1). When individuals make decisions based on popular information and not credible information, they are not engaging in elaborate reasoning (i.e. taking the time to think through what they've read). ELM considers that information seekers make judgments in two ways: (1) fast with simple reasoning and (2) slow with elaborate reasoning (Petty Richard and Cacioppo, 1986).

Since social media sites provide information in a user-friendly interface, information seekers are continuously tempted to make fast judgments based on limited information. Additionally, the ELM model doesn't provide ways to characterize the interplay between information dimensions and types of decisions. Social media sites provide content moderation methods that help users curate their information feed, which can aid in finding relevant scientific information. This can also aid in supporting elaborate thinking. However, if users deem these methods unreliable, they will not engage in content moderation increasing the likelihood of information overload. The Intent Model can aid in characterizing what types of information people seek during a crisis for different types of decisions (Figure 1.3).

Using our intent model, an information seeker can triage information based on different intent categories. For each intent category, additional details are accessible to the user via the source, content, and attitude or the messages. The message's source is defined as an expert or general user and the degree of subjectivity. The content and attitude refer to the topic and sentiment of the message. Most content analysis studies generated a labeling guide based on previous literature reviews, such as guiding principles for classifying social media news articles (MacKay et al., 2021), informativeness (whether a tweet contains relevant information or not) (Olteanu et al., 2015), an existing crisis communication model like HBM (Myneni et al., 2023). We chose to follow this approach for our labeling guide. During information-seeking behavioral patterns, consumers pay attention to the message source, So two graduate students took into account the source of the message (expert/general user) and the message content to determine the message's intent. We classified 6,844 tweets into five intent categories: (i) Expert Guidance, (ii) Situational Awareness (iii) Severity (iv) Reactions, (v) Lived Experience. Figure 2 provides a topic distribution and an example tweet in each intent category with similar uncertainty profiles to emphasize the differences across each intent category.

Expert Guidance categorizes tweets from experts that are providing some suggestions or recommendations to address the pandemic (21% of sample) (Wang et al., 2021; Brady et al., 2023; Ehrmann

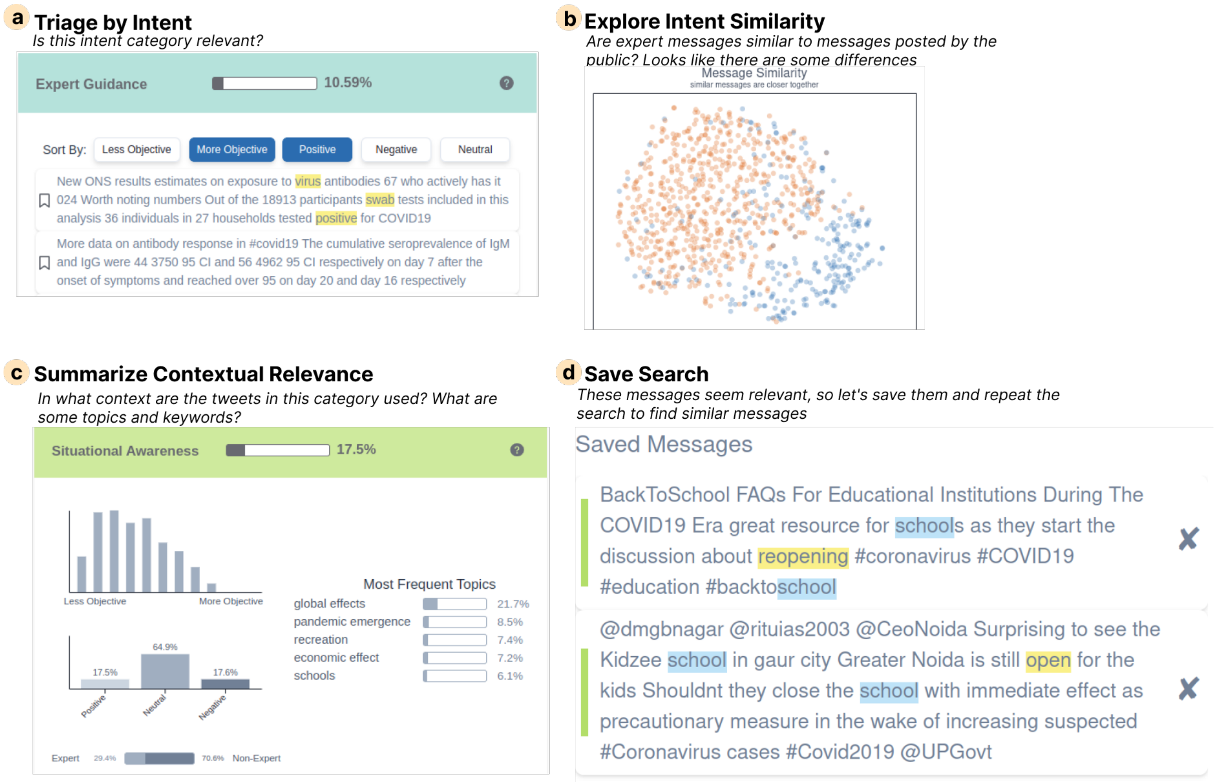

Figure 3: A hypothetical workflow of a social media user leveraging the intent classification model for sensemaking during the pandemic.

and Wabitsch, 2022). The most frequent topics for this category are scientific information, testing, and vaccine development. *Situational Awareness* are tweets from any source that are providing updates on news alerts, business proceedings, and policy revisions (33% of sample) (MacKay et al., 2021; Myneni et al., 2023). The most frequent topics for this category are global effects, pandemic emergence, and recreation. *Severity* describes tweets that qualitatively or quantitatively report on the impact of the pandemic (11% of sample) (Myneni et al., 2023). The most frequent topics for this category are pandemic emergence, global effects, and case statistics. *Reactions* are tweets from any source that describe emotions, comments, or responses towards events caused by the pandemic (30% of sample) (Wang et al., 2021). The most frequent topics for this category are testing, recreation, and global effects. *Lived Experience* tweets describe direct personal experiences with specific details regarding location and time (8% of sample) (Wang et al., 2021). The most frequent topics for this category are global effects, schools, and recreation.

A DistilBERT (Sanh et al., 2019) **classification model** was trained on the 6844 labeled tweets until overall accuracy reached approximately 70%, indicating good performance. The accuracy per intent category: (1) Expert Guidance 69%, (2) Reactions 74% (3) Lived Experience 80% (4) Severity 93% (5) Situational Awareness 74%. These results on real-world data indicated that our model could also be used to learn from user interactions. Current content moderation methods do not allow users to update the underlying model. However, for machine-guided social media systems to better address the changing needs of information consumers, users need to update the underlying models to match their mental model of changing information. The Intent Model allows users to update each message's intent to one that aligns more with the user's mental model or curate their intent category based on the source, content, and attitude of the messages.

5 Sensemaking during a Crisis

We present an AI-driven sensemkaing scenario during the COVID-19 pandemic to demonstrate the utility of our proposed intent classification model (ICM). We integrated ICM (a predictive model trained on our collection of tweets) with a web-based interface, which allows lay information consumers to triage credible and potentially actionable information. Let us see how this interface can be used by Mary, a parent whose child currently studies at home due to COVID-19-affected school closing. She needs to decide if she should reduce her working hours and invest in homeschooling. Figure 3 shows Mary's workflow as she interacts with our interface. From her experience, she knows that not all tweets are reliable and that topics provided by Twitter do not give her sufficient control to navigate the relevant messages. She uses the interface to triage the relevant information using *Situational Awareness* as the intent category of interest (Figure 3.a). Additionally, she selects all intent categories to learn more about the contrasting messages by reviewing the representative tweets in each category. She finds that by using the intent categories and the highlighted words she can gradually make sense of the different facets of the conversation related to the pandemic that might or might not be related to her decision-making goal. She can easily filter out messages from categories she is not interested in such as Reactions and Severity.

To find more contextually relevant messages, Mary filters by keyword *school* and chooses to summarize the overall message profile and observes the patterns across sentiment, subjectivity, expertise, and topics across each intent category (Figure 3.c). She notices that Situational Awareness tweets are from experts and lay users and that *schools* is a frequent topic. She notices that messages from experts and general users have distinct groupings which tell her they may have different opinions (Figure 3.b). Interestingly, she finds the helpful messages discussing distance learning and back-to-school efforts come from lay users because of their lived experience. It did not occur to her that online learning could be a possible solution for her child to continue learning. She found tweets suggesting reopening in-person learning once proper safety measures are determined. She saves those messages and repeats her search to find even more similar tweets (Figure 3.d).

She also decided to filter the selection criteria only to consider tweets from the *schools* topic. As she expected, applying this filter allowed her to see more messages about other parents' experiences with both homeschooling and online learning. After reviewing her saved tweets, she decided that online learning is a viable option. Since Mary has a day job she needs to report to, online learning for her child would allow Mary to continue working and her child to continue receiving an education. By leveraging the intent model, Mary could quickly reduce her exposure to irrelevant tweets. She could also assess the credibility of the messages critically and have sufficient control over the information she needed to know.

6 Conclusion and Future Work

In this paper, we present a crisis-related intent classification model and present its utility via examples and usage scenarios using the COVID-19 pandemic as an example. We developed and trained classification models for sentiment, subjectivity, and topic to further our understanding of how experts and general users communicate during the initial stages of a crisis. We use Twitter messages as the basis of our analysis to profile the information uncertainty and address the need for principled approaches towards sensemaking of socially mediated information during a crisis. We are currently developing a software prototype, demonstrated in Section 5, that allows lay users to explore messages using the intent of a message author and control their exposure to crisis-relevant information, by focusing on what they would need to address their pressing questions.

Future work would consider a user study with diverse participants to understand how users perceive the usefulness and utility of our proposed sensemaking workflow and the resulting human-machine interface in their information-seeking processes. We intend to build upon and further develop the web interface, taking into account previous studies on end-user content moderation techniques.

References

Alaa Abd-Alrazaq, Dari Alhuwail, Mowafa Househ, Mounir Hamdi, and Zubair Shah. 2020. Top concerns of tweeters during the covid-19 pandemic: infoveillance study. *Journal of medical Internet research*, 22(4):e19016.

Benigno E Aguirre and Kathleen J Tierney. 2001. Testing shibutani's prediction of information seeking behavior in rumor.

Sakun Boon-Itt, Yukolpat Skunkan, et al. 2020. Public perception of the covid-19 pandemic on twitter: sentiment analysis and topic modeling study. *JMIR Public Health and Surveillance*, 6(4):e21978.

William J Brady, Joshua Conrad Jackson, Björn Lindström, and MJ Crockett. 2023. Algorithm-mediated social learning in online social networks. *Preprint at OSF preprints. https://doi. org/10.31219/osf. io/yw5ah*.

Umar Ali Bukar, Marzanah A Jabar, Fatimah Sidi, Rozi Nor Haizan Nor, Salfarina Abdullah, and Iskandar Ishak. 2022a. Mobile social media use in flooding emergencies: The impact of information seeking and sharing through social media crisis communication and resilience model. In *2022 Applied Informatics International Conference (AiIC)*, pages 51–56. IEEE.

Umar Ali Bukar, Fatimah Sidi, Marzanah A Jabar, Rozi Nor Haizan Binti Nor, Salfarina Abdullah, and Iskandar Ishak. 2022b. A multistage analysis of predicting public resilience of impactful social media crisis communication in flooding emergencies. *IEEE Access*, 10:57266–57282.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

LaShonda Louallen Eaddy and Yan Jin. 2018. Crisis history tellers matter: The effects of crisis history and crisis information source on publics' cognitive and affective responses to organizational crisis. *Corporate Communications: An International Journal*.

Elemental Editors. 2020. Six months in 50 experts to trust in a pandemic.

Michael Ehrmann and Alena Wabitsch. 2022. Central bank communication with non-experts–a road to nowhere? *Journal of Monetary Economics*, 127:69–85.

Tarleton Gillespie. 2022. Do not recommend? reduction as a form of content moderation. *Social Media+ Society*, 8(3):20563051221117552.

Melissa W Graham, Elizabeth J Avery, and Sejin Park. 2015. The role of social media in local government crisis communications. *Public Relations Review*, 41(3):386–394.

Maarten Grootendorst. 2020. Bertopic: Leveraging bert and c-tf-idf to create easily interpretable topics.

Andreas Hanselowski, Christian Stab, Claudia Schulz, Zile Li, and Iryna Gurevych. 2019. A richly annotated corpus for different tasks in automated fact-checking. *arXiv preprint arXiv:1911.01214*.

Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claimbuster. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1803–1812.

Md Arafat Hossain, Steve Versteeg, Jun Han, Muhammad Ashad Kabir, Jiaojiao Jiang, and Jean-Guy Schneider. 2018. Mining accurate message formats for service apis. In *2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, pages 266–276. IEEE.

Khairul Islam, Najma Akhther, and Matthew W Seeger. 2023. Variability in media content of public health outbreak coverage: A crisis communication approach. *Communication Studies*, 74(2):113–130.

Kyungeun Jang and Young Min Baek. 2019. When information from public health officials is untrustworthy: The use of online news, interpersonal networks, and social media during the mers outbreak in south korea. *Health communication*, 34(9):991–998.

Shagun Jhaver, Alice Qian Zhang, Quan Ze Chen, Nikhila Natarajan, Ruotong Wang, and Amy X Zhang. 2023. Personalizing content moderation on social media: User perspectives on moderation choices, interface design, and labor. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW2):1–33.

Hye Kyung Kim and Sung-Un Yang. 2009. Cognitive processing of crisis communication: Effects of csr and crisis response strategies on stakeholder perceptions of a racial crisis dynamics. *Public Relations Journal*, 3(1):1–39.

Xialing Lin, Patric R Spence, Timothy L Sellnow, and Kenneth A Lachlan. 2016. Crisis communication, learning and responding: Best practices in social media. *Computers in human behavior*, 65:601–605.

Brooke Fisher Liu, Julia Daisy Fraustino, and Yan Jin. 2016. Social media use during disasters: How information form and source influence intended behavioral responses. *Communication Research*, 43(5):626–646.

Brooke Fisher Liu, Yan Jin, Lucinda Austin, Erica Kuligowski, and Camila Espina Young. 2020. 14 the social-mediated crisis communication (smcc) model. *Advancing Crisis Communication Effectiveness: Integrating Public Relations Scholarship with Practice*, page 63.

Joanne Chen Lyu, Eileen Le Han, and Garving K Luli. 2021. Covid-19 vaccine–related discussion on twitter: topic modeling and sentiment analysis. *Journal of medical Internet research*, 23(6):e24435.

Melissa MacKay, Taylor Colangeli, Daniel Gillis, Jennifer McWhirter, and Andrew Papadopoulos. 2021. Examining social media crisis communication during early covid-19 from public health and news media for quality, content, and corresponding public sentiment. *International Journal of Environmental Research and Public Health*, 18(15):7986.

Charles W Meadows, Cui Zhang Meadows, Lu Tang, and Wenlin Liu. 2019. Unraveling public health crises across stages: Understanding twitter emotions and message types during the california measles outbreak. *Communication Studies*, 70(4):453–469.

Sebastião Miranda, David Nogueira, Afonso Mendes, Andreas Vlachos, Andrew Secker, Rebecca Garrett, Jeff Mitchel, and Zita Marinho. 2019. Automated fact checking in the news room. In *The World Wide Web Conference*, pages 3579–3583.

Corrado Monti, Luca Maria Aiello, Gianmarco De Francisci Morales, and Francesco Bonchi. 2022. The language of opinion change on social media under the lens of communicative action. *Scientific Reports*, 12(1):17920.

Sahiti Myneni, Paula Cuccaro, Sarah Montgomery, Vivek Pakanati, Jinni Tang, Tavleen Singh, Olivia Dominguez, Trevor Cohen, Belinda Reininger, Lara S Savas, et al. 2023. Lessons learned from interdisciplinary efforts to combat covid-19 misinformation: Development of agile integrative methods from behavioral science, data science, and implementation science. *JMIR infodemiology*, 3(1):e40156.

Preslav Nakov, David Corney, Maram Hasanain, Firoj Alam, Tamer Elsayed, Alberto Barrón-Cedeño, Paolo Papotti, Shaden Shaar, and Giovanni Da San Martino. 2021. Automated fact-checking for assisting human fact-checkers. *arXiv preprint arXiv:2103.07769*.

Usman Naseem, Imran Razzak, Matloob Khushi, Peter W Eklund, and Jinman Kim. 2021. Covidsenti: A large-scale benchmark twitter data set for covid-19 sentiment analysis. *IEEE transactions on computational social systems*, 8(4):1003–1015.

Francisco Braulio Oliveira, Amanul Haque, Davoud Mougouei, Simon Evans, Jaime Simao Sichman, and Munindar P Singh. 2022. Investigating the emotional response to covid-19 news on twitter: a topic modeling and emotion classification approach. *IEEE Access*, 10:16883–16897.

Alexandra Olteanu, Sarah Vieweg, and Carlos Castillo. 2015. What to expect when the unexpected happens: Social media communications across crises. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*, pages 994–1009.

E Petty Richard and John T Cacioppo. 1986. Communication and persuasion: Central and peripheral routes to attitude change.

Juan Manuel Pérez, Juan Carlos Giudici, and Franco Luque. 2021. pysentimiento: A python toolkit for sentiment analysis and socialnlp tasks.

Christian Reuter and Marc-André Kaufhold. 2018. Fifteen years of social media in emergencies: a retrospective review and future directions for crisis informatics. *Journal of contingencies and crisis management*, 26(1):41–57.

Constanze Rossmann, Lisa Meyer, and Peter J Schulz. 2018. The mediated amplification of a crisis: Communicating the a/h1n1 pandemic in press releases and press coverage in europe. *Risk analysis*, 38(2):357–375.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Stefan Stieglitz, Milad Mirbabaie, Lara Schwenner, Julian Marx, Janina Lehr, and Felix Brünker. 2017. Sensemaking and communication roles in social media crisis communication.

Sarah E Vaala, Matthew B Ritter, and Deepak Palakshappa. 2022. Experimental effects of tweets encouraging social distancing: Effects of source, emotional appeal, and political ideology on emotion, threat, and efficacy. *Journal of Public Health Management and Practice*, 28(2):E586–E594.

Ranjit Vijayan. 2021. Teaching and learning during the covid-19 pandemic: A topic modeling study. *Education Sciences*, 11(7):347.

Yan Wang, Haiyan Hao, and Lisa Sundahl Platt. 2021. Examining risk and crisis communications of government agencies and stakeholders during early-stages of covid-19 on twitter. *Computers in human behavior*, 114:106568.

Lisa Washburn. 2020. Understanding the health belief model. *University of Tennessee: Knoxville, TN, USA*.