

Can LLMs Assist Annotators in Identifying Morality Frames? - Case Study on Vaccination Debate on Social Media

Tunazzina Islam
Department of Computer Science,
Purdue University
West Lafayette, IN-47907, USA
islam32@purdue.edu

Dan Goldwasser
Department of Computer Science,
Purdue University
West Lafayette, IN-47907, USA
dgoldwas@purdue.edu

Abstract

Nowadays, social media is pivotal in shaping public discourse, especially on polarizing issues like vaccination, where diverse moral perspectives influence individual opinions. In NLP, data scarcity and complexity of psycholinguistic tasks, such as identifying morality frames, make relying solely on human annotators costly, time-consuming, and prone to inconsistency due to cognitive load. To address these issues, we leverage large language models (LLMs), which are adept at adapting new tasks through few-shot learning, utilizing a handful of in-context examples coupled with explanations that connect examples to task principles. Our research explores LLMs' potential to assist human annotators in identifying morality frames within vaccination debates on social media. We employ a two-step process: generating concepts and explanations with LLMs, followed by human evaluation using a "think-aloud" tool. Our study shows that integrating LLMs into the annotation process enhances accuracy, reduces task difficulty, lowers cognitive load, suggesting a promising avenue for human-AI collaboration in complex psycholinguistic tasks.

CCS Concepts

• Computing methodologies → Natural language processing;
• Information systems → Social networks; • Human-centered computing → Empirical studies in collaborative and social computing.

Keywords

social media, large language models, natural language processing, vaccine, moral foundation, COVID-19, morality frames

ACM Reference Format:

Tunazzina Islam and Dan Goldwasser. 2025. Can LLMs Assist Annotators in Identifying Morality Frames? - Case Study on Vaccination Debate on Social Media. In *Proceedings of the 17th ACM Web Science Conference 2025 (WebSci '25)*, May 20–24, 2025, New Brunswick, NJ, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3717867.3717902>

1 Introduction

The COVID-19 pandemic marked an unprecedented moment in global history, not only as a public health crisis but also as a pivotal event in the use of digital platforms. Social media has significantly

influenced public discourse, particularly in the context of highly polarized topics such as vaccinations [2, 9, 22, 27, 30, 38, 55, 61, 67, 68, 71] by playing a dual role in both disseminating vital information [28, 79–81] and spreading misinformation [20, 63, 74]. This *infodemic* [71] has complicated public health efforts, particularly around contentious issues like vaccination, where moral and ethical concerns often influence individual and collective behaviors.

Debates about vaccination are often influenced by diverse moral perspectives [3, 6, 30, 64, 76]. Understanding how these moral perspectives shape public discourse is crucial for designing better tools and interventions in online environments. Morality frame analysis plays a key role in vaccine debates, as it helps to understand how moral beliefs influence individual and public opinions on vaccination. It provides a comprehensive framework for capturing the conveyed moral sentiments by identifying the relevant moral foundation (MF) [24–26] and at a fine-grained level, the moral sentiment expressed towards the entities involved [57, 66]. Moral Foundation Theory (MFT) [25, 26] suggests a theoretical framework for analyzing six moral values (i.e., foundations, each with a positive and a negative polarity) central to human moral sentiment (Table 1). Let's look into the example text and morality frames:

Example 1: *Pfizer vaccine testing utilized cell lines from human fetus tissue. This makes it abhorrent, heretical and blasphemous to anyone calling themselves a Christian.*

Analysis: Morality frame analysis identifies the moral attitudes expressed in the text (degradation), and how different entities mentioned in it are perceived ("Pfizer vaccine" is degrading, "Christian" feels disgusted).

Our morality analysis is driven by social science research that highlights the connection between moral foundation preferences and COVID-19 vaccine-related health choices [13, 17, 60]. Pagliaro et al. [60] have demonstrated that the endorsement of moral principles such as *fairness* and *care*, as opposed to *loyalty* and *authority*, was found to correlate positively with *trust in science*. For fine-grained morality frame analysis, we employ the recently proposed morality-frame formalism [57, 66] that identifies moral roles associated with moral foundation expressions in text. These roles correspond to **actor/target** roles and **positive** or **negative** polarity, which can be interpreted in the context of a specific moral foundation. In Example 1, "Pfizer vaccine" is the **negative actor** in the context of degradation which causes disgust among "Christian" (**negative target**) and the reason stems from the "use of fetal cell lines in vaccine development, framing this as a violation of sacred values, specifically within the Christian context".



CARE/HARM: : Indicates that an individual other than the speaker deserves care or suffers harm. Rooted in the virtues of kindness, gentleness, and nurturance.
FAIRNESS/CHEATING: Principles of justice, individual rights, and autonomy; comparisons with other groups. Equal opportunities. Social opposition to “Free Riders”.
LOYALTY/BETRAYAL: Founded on the virtues of patriotism and self-sacrifice for the group. Engaged whenever there is a sentiment of “one for all, and all for one.”
AUTHORITY/SUBVERSION: Involves deference (or opposition) to recognized authority and adherence to traditions. It encompasses social order and the responsibilities inherent in hierarchical relationships, including obedience, respect, and the execution of role-specific duties.
SANCTITY/DEGRADATION: More than a religious value, it emphasizes respect for the human spirit and societal distaste for personal degradation, which typically evokes disgust. This principle concerns both physical and spiritual contagion and promotes virtues such as chastity, wholesomeness, and restraint of desires. It reinforces the common belief that the body is a temple susceptible to desecration through immoral actions and contaminants, a notion not exclusive to religious traditions.
LIBERTY/OPPRESSION: Reflects individuals’ reactions to those who dominate and infringe upon their freedom. Disdain for bullies and authoritarian figures often inspires unity and solidarity among people, motivating them to resist or overthrow their oppressors.

Table 1: Six basic moral foundations [25, 26].

Text: Surprise: Fox News Hosts Are Following Strict COVID Protocols While Telling Viewers Masks and Vaccines Are Liberal Plots.
Question1: What’s the moral foundation of the above text? Give the explanation of choosing the certain answer.
Answer1: fairness/cheating.
Explanation1: because the text implies that Fox News hosts are being hypocritical by following strict COVID protocols while telling their viewers that masks and vaccines are liberal plots, which is unfair.
Question2: Answer the actor-target role with polarity of corresponding moral foundation. Please only choose actor-target role with polarity from the given text. Don’t generate a new actor-target role with polarity. Give the explanation of choosing the certain answer.
Answer2: Negative Actor: Fox News hosts. Negative Target: Viewers.
Explanation2: Fox News hosts are the negative actors because they are being hypocritical, and the viewers are the negative targets because they are being misled by the hosts’ contradictory actions and statements.

Figure 1: Morality frame identification with explanation. Inputs are shown in blue, and outputs are shown in red.

Given the complexity and the psycholinguistic nuances embedded in contentious topics, analyzing morality frames poses a substantial challenge, especially in the field of NLP. This challenge arises primarily from three factors. **Firstly**, the nuances can manifest in very different ways depending on context. For example, the concept of **loyalty** can be discussed in the context of the climate debate (e.g., “support American energy”) or in the context of vaccination (e.g., “take COVID-19 vaccine to protect your community and friends”). **Secondly**, building models using supervised learning techniques requires extensive labeled datasets, which can be resource-intensive to create. **Thirdly**, interpreting psycholinguistic cues demands specialized knowledge and places a cognitive load on human annotators. Recent advancements in Artificial Intelligence (AI), particularly the development of Large Language Models (LLMs), present new opportunities to augment human capabilities in this domain. LLMs have demonstrated the ability to perform complex tasks with minimal input through few-shot learning [11], where the model is given a few examples to learn a new task. This capability can be leveraged to assist human annotators by generating initial labels and explanations for moral frames, thereby reducing their cognitive burden and improving the consistency and accuracy of annotations. In this paper, we investigate whether the newly emerged paradigm in NLP; few-shot prompting of LLMs [11] and the practice of providing explanations of answers when learning from examples in-context [43] —is better equipped to address those challenges.

Explanations are fundamental to human learning [1], as they underscore task principles that facilitate broad generalizations [46, 47]. Consider the example text (“*Surprise: Fox News Hosts Are Following Strict*

COVID Protocols While Telling Viewers Masks and Vaccines Are Liberal Plots.”) from Fig. 1 for morality frame identification task. An explanation can elaborate on a brief answer (e.g., **fairness/cheating**) by connecting it to the broader reasoning process necessary to solve the problem (e.g., “the text implies that Fox News hosts are being hypocritical by following strict COVID protocols while telling their viewers that masks and vaccines are liberal plots, which is unfair”). Thus, explanations enhance understanding by demonstrating how task principles connect questions to their answers. Unlike previous methods of morality frame identification [57, 65, 66] that relied on annotated data [57, 66] or solely zero/few-shot prompting [65], our approach uses few-shot prompting with explanations through in-context learning with LLMs [11, 16, 40, 44]. We describe the methodology in section 3.

We have designed an evaluation task for LLMs-generated morality frames that utilizes a dedicated web application¹. This platform also provides explanations generated by the LLMs. If annotators believe the predictions are incorrect, they can submit the correct answers. The goal is to determine whether the LLMs’ outputs accurately align with the expected morality frame expressed in the text, thereby facilitating a systematic assessment of the model’s performance in handling complex psycholinguistic tasks. Additionally, we conduct a survey to gather feedback from participants about their experiences. We ask whether the explanations are helpful and how they improve the process, if the explanations reduce cognitive load and in what ways, the difficulty of the task rated on a scale of 1 to 5, and the time required to complete each batch of annotations. We detail the task in subsection 3.2. In this paper, we offer three primary contributions:

Innovative Framework for Human-AI Collaborative Annotation: We introduce a novel framework that combines the strengths of LLMs with human expertise to enhance the process of morality frame identification in social media discourse. By leveraging few-shot learning with explanations, our approach not only improves the accuracy and efficiency of annotations but also significantly reduces the cognitive load on human annotators. This work demonstrates how AI can be integrated into complex psycholinguistic tasks.

¹Our tool is available at <https://github.com/tunazislam/llms-annotator>

Development of Web-based Annotation Tool: We have developed a specialized web-based tool that facilitates a collaborative annotation process, enabling annotators to interact with AI-generated labels and explanations. This tool is designed to support human decision-making by providing clear, context-sensitive guidance, thereby improving the overall quality of the annotation process. Our tool helps evaluate how effectively LLMs identify and interpret morality frames within complex texts.

Comprehensive Empirical Evaluation and Insights: Through an extensive empirical study, we evaluate the effectiveness of LLMs in assisting human annotators, focusing on key metrics such as accuracy, task difficulty, cognitive load, and time efficiency. Our findings reveal that LLMs can serve as effective collaborators, offering substantial improvements in task performance while maintaining the critical role of human judgment.

2 Related Works

The rise of social media has drastically transformed public discourse, particularly on contentious topics [4, 30, 32, 34, 35, 53, 61, 69, 70], by providing a platform that can both inform [7, 12, 15] and mislead [52, 62, 72, 77] on a massive scale. This transformation has been vividly illustrated during the COVID-19 pandemic, which sparked a global infodemic [71]. Ferrara et al. [19] studied the narratives fueled with conspiracy theories and COVID-19 misinformation on the global news sentiment, on hate speech and social bot interference, and on multimodal Chinese propaganda. Wang et al. [74] developed a codebook to characterize the various types of COVID-19 misinformation images related to the virus, from false medical advice to conspiracy theories and analyzed how COVID-19 misinformation images are used on social media.

Social media platforms have become battlefields where moral perspectives clash, underscoring the necessity for a nuanced analysis of how moral beliefs shape public opinion and individual positions in vaccine debates [3, 6, 30, 64, 76]. MFT [25, 26] suggests a theoretical framework for analyzing the moral perspective, containing six basic moral foundations (Table 1). Weinzierl and Harabagiu [76] identified vaccine hesitancy profiles in COVID-19 discussions by analyzing the underlying MFs. Islam and Goldwasser [30] analyzed COVID-19 vaccine campaigns on Facebook by predicting MF and themes. Beiró et al. [6] employed the MFT to assess the moral narratives around vaccination debate on Facebook.

Recently, Roy et al. [66] introduced a framework for analyzing moral sentiment called morality frames building on and extending MFT. Pacheco et al. [57] adapted this morality frame formalism that identifies moral roles associated with moral foundation expressions in the text. Pacheco et al. [57], Roy et al. [66] used annotated data to train a relational classifier using a declarative framework for specifying deep relational models called DRaIL [56]. One limitation of these works is their dependency on labeled datasets annotated by humans. In contrast, we rely on LLMs inference to generate labels and explanations for morality frames in this paper.

Recent studies show that LLMs can perform tasks with just a few examples in-context [11, 16, 40, 44]. A variety of prior work has explored task instructions as a part of few-shot prompt [51] and can benefit from explicitly decomposing problems into multiple steps [50]. Other recent work has shown that LLMs can benefit

Instructions	Instructions: 1. Answer moral foundation of given text. If the text does not make any moral judgement, answer "none". For this case, you don't have to answer actor-target polarity. 2. If you answer a moral foundation other than "none", you need to answer actor-target polarity. 3. Answer the actor-target role with polarity of corresponding moral foundation. Answer who the Positive Actor / Positive Target / Negative Actor / Negative Target is. 4. Provide explanations of your answers.
Definition of Moral Foundations	For answering Moral Foundation, please look into the following categories: 1. care/harm -> "definition" 2. fairness/justice -> "definition" 3. loyalty/authority -> "definition" 4. authority/subversion -> "definition" 5. sanctity/degradation -> "definition" 6. liberty/oppression -> "definition" 7. none -> "definition"
Definition of Roles	For Answering actor and target roles with polarity, please look into following definitions: "definition" Text: "Text" Question: "What is the moral foundation of the above text? Give the explanation of choosing the certain answer." Answer: "moral foundation"
7 examples with explanations	Question: "Answer the actor-target role with polarity of corresponding moral foundation. Please only choose actor-target role with polarity from the given text. Don't generate a new actor-target role with polarity. Give the explanation of choosing the certain answer." Answer: "moral foundation polarity role" Explanation: "role explanation"
Test Text	Text: "The text" Question: "What is the moral foundation of the above text? Give the explanation of choosing the certain answer." Answer: "moral foundation" Question: "Answer the actor-target role with polarity of corresponding moral foundation. Please only choose actor-target role with polarity from the given text. Don't generate a new actor-target role with polarity. Give the explanation of choosing the certain answer." Answer: "moral foundation polarity role" Explanation: "role explanation"

Figure 2: Prompt template for morality frame identification with explanation. Blue colored segment is input prompt and red colored segment is generated output by LLMs.

from examples that decompose the reasoning process (can be seen as an explanation) leading to an answer [75]. Lampinen et al. [43] provided explanations after the answer in the prompt for few-shot settings.

In recent times, researchers have been exploring assistive possibilities of LLMs in several tasks [39, 41, 45, 78]. Kolla et al. [41] examined the feasibility of using LLMs for online moderation, i.e., identifying rule violations on Reddit. Lin et al. [45] proposed LLM-assisted macro revisions. However, researchers have evaluated the performance of LLMs as annotators in various NLP tasks [18, 73]. Mei et al. [48] have demonstrated that LLMs can serve effectively as an assistive tool for annotation tasks. Gilardi et al. [23], Huang et al. [29] have investigated LLM's ability to independently perform annotations, showcasing its potential to enhance both the efficiency and accuracy of these processes.

In contrast to previous works, our work explores the few-shot prompting with explanations capacity of LLMs to assist annotators in complex psycholinguistic tasks, specifically within the context of identifying morality frames in vaccination debates on social media. Utilizing a two-step method, we first generate labels and explanations via few-shot prompting with LLMs, followed by a human evaluation phase that includes a "think-aloud" tool to provide judgments and collect feedback.

3 Methodology

In this section, we describe our collaborative annotation framework that leverages the strengths of LLMs to enhance the accuracy and efficiency of human annotations in identifying morality frames within social media discourse.

3.1 Identification and Explanation of Morality Frames with LLMs

We approach the task of identifying moral foundations as a text-generation problem, where the model is prompted in a few-shot manner to generate both the moral foundation label for a given text and provide the reasoning behind this choice. We engineer prompts in a few-shot style to generate the actor-target role with the polarity of the corresponding moral foundation and to provide an explanation. For each moral foundation of a given text, we prompt the LLMs to predict two general role types: **actor** and **target**, each with an associated polarity (*positive*, *negative*). An

Text: #BidenIsADisgrace can't back up his rhetoric regarding Mr. Trump's vaccine stances and mandate a #COVID19 vaccine. He'd rather cheerlead for businesses, schools and any other groups to do his work for him. Why can't we have a President that actually stands up?? #BidenIsAFailure

Moral Foundation: authority/subversion
MF Explanation: because the text criticizes President Biden's actions and leadership, implying that he is not doing his job properly and should be taking a stronger stance
Actor-Target polarity role: positive actor: , negative actor: Biden, positive target: , negative target: Americans (implied)
Role Explanation: President Biden is the negative actor who is not taking proper action, and Americans are the negative targets who are affected by his inaction

Do you think the given moral foundation and actor-target polarity role of the provided text are correct?
☐ yes ☐ no

After finishing the task, please click Submit button.

Showing Task 1 of 11

(a) Task interface in general.

Text: This IS illegal- Forcing people against their will, religion, beliefs and giving injections is a violation our AMERICAN CITIZEN'S RIGHTS!!! You can't legally inject humans when this shot! It is NOT a vaccine, falsely labeling the injection as a vaccine is ILLEGAL! Law suits!

Moral Foundation: liberty/oppression
MF Explanation: because the text expresses resentment toward those forcing others to go against their will and rights, which the author perceives as oppression
Actor-Target polarity role: positive actor: , negative actor: those who are forcing injections, positive target: , negative target: people, AMERICAN CITIZEN'S RIGHTS
Role Explanation: The "people" and "AMERICAN CITIZEN'S RIGHTS" are being negatively impacted by the forcible actions of others (those who are forcing injections)

Do you think the given moral foundation and actor-target polarity role of the provided text are correct?
☒ yes ☐ no

After finishing the task, please click Submit button.

(b) Task interface (annotators agree with LLMs generated moral foundation and role).

Text: When I see **people** rejecting the vaccine in the south, I don't only see Trump supporters or conspiracy nutters. What I see is a lot of fear and distrust, **that** must have inspired this heightened paranoia of a violent world and medical establishment.

Moral Foundation: care/harm
MF Explanation: because the author is expressing concern for people who are fearful and distrustful due to past hurts and betrayals
Actor-Target polarity role: positive actor: , negative actor: past hurt and betrayal, positive target: , negative target: People rejecting the vaccine/Trump supporters/conspiracy nutters
Role Explanation: Because the text implies that the past hurt and betrayal (Negative Actor) have caused fear, distrust, and paranoia (Negative effect) in the people rejecting the vaccine/Trump supporters/conspiracy nutters (Negative Target)

Do you think the given moral foundation and actor-target polarity role of the provided text are correct?
☐ yes ☒ no

If you feel that the given moral foundation is not correct, choose the correct moral foundation from following:
☐ care/harm ☐ interest/reason ☒ liberty/justice ☐ authority/subversion ☐ security/legislation ☐ liberty/oppression ☐ none

First pick the color.
 Second highlight the text for actor-target role with polarity associated with corresponding moral foundation.

(c) Task interface (error in moral foundation and role).

Figure 5: Task interface for assessing LLMs generated morality frames.

along with explanations. The main task interface for annotation is shown in Fig. 5.

3.3 Participants' Experiences Survey

We conduct a survey to collect feedback from participants who interacted with our system. The survey aims to understand the annotators' experiences, focusing on several key areas:

Task Difficulty: Annotators evaluate the difficulty of the task on a scale from 1 to 5 where 1 indicates "very easy", 2 indicates "easy", 3 indicates "okay", 4 indicates "hard", and 5 indicates "very hard". They are asked how difficult the morality frame identification task is with and without labels and explanations.

Effectiveness of Explanations: Annotators are asked to answer whether the explanations provided by the LLMs are helpful in understanding the morality frames within the texts they analyze.

Cognitive Load: We ask annotators if the explanations help reduce the cognitive loads, making it easier for them to perform their tasks without feeling overwhelmed by the complexity of the content.

Time Efficiency: Annotators report the time required to complete each batch of annotation. This measure helps gauge the efficiency of the LLMs in aiding annotators and whether the tool streamlines their workflow.

Comments on Explanation Efficiency: Annotators are encouraged to offer detailed comments on the specific ways in which the provided explanations facilitate a reduction in task complexity and ease cognitive demands, thereby enhancing their overall annotation experience.

4 Experimental Details

In this paper, we conduct our study on a timely topic, the COVID-19 vaccine debate, centered around discussions on social media. We utilize the dataset introduced by Pacheco et al. [57], which consists of 750 tweets written in English about the COVID-19 vaccine from April to October 2021. Our task focuses **solely on the tweet text** and **does not** consider other attributes from that dataset.

4.1 Few-Shot Learning with Explanation

One of our goals is to leverage LLMs' few-shot prompting with explanation capabilities to identify morality frames discussed in a text. We use GPT-4o² [54] for our experiment. At the beginning of the prompt, we provide an overview of the task along with a description of the expected labels as part of the task instructions before introducing the few-shot examples. We show the prompt template in Fig. 2. To alleviate the issue of being hallucinated, we force LLMs to answer the questions by selecting from the given moral foundation category only. For each task, we construct 7-shot prompts that include examples with explanations covering all six moral foundations (corresponding actor-target role with polarity) and non-moral ("none") cases. If the moral foundation of the text is "none", the corresponding actor-target-polarity will be "none". We place the explanation on a line after the answer (both for moral foundation and actor-target-polarity), preceded by "Explanation:" [43].

4.2 Human Evaluation

Our specialized web tool (Fig. 5) allows annotators to assess moral foundations (with explanation) and actor-target polarity roles (with explanation) generated by LLMs. If they believe that the LLMs-generated answer is correct, they can select 'yes' and move on to the next (Fig. 5b). If they believe that the LLMs-generated answer is not correct, they can select a specific moral foundation and highlight the corresponding actor-target roles with polarity in the provided text (Fig. 5c). The graphical interface provides necessary instructions (Fig. 3), eight examples with explanations (Fig. 4), and two practice examples.

4.2.1 Participants. We compile a list of potential participants in preparation for the evaluation of our task. The prerequisites for inclusion on this list are that candidates must have at least an undergraduate-level academic degree. The initial selection is based on our academic networks. We have recruited nine ($N = 9$) participants—five males and four females, aged between 25 and 50. The group includes graduate students (both Ph.D. and Masters), a postdoctoral researcher, and a faculty member. Five of the nine participants are researchers in the NLP and computational social science (CSS) field. One of the participants is a Software Engineer. One participant works in Machine Learning (ML) in the Medical

²<https://openai.com/index/hello-gpt-4o/>

Surgery domain. Another is a Security researcher. Another participant has just finished their undergraduate degree and started their Masters, doing research in ML and Software Engineering. All participants reside in the United States.

We have developed internal guidelines that set ethical and legal standards for our research group. These guidelines include sending all potential participants a thorough information upon first contacting them, which allows them to determine whether or not they are willing to participate in the study. We inform them that there will be no compensation for their involvement. Upon agreeing to participate, they receive another detailed information sheet that explains the study’s background, outlines the process and evaluation methods, and describes how we will use and store the collected data. All interviews have been conducted remotely using the video conferencing software ‘Zoom’.

When we contact the participants, we inform them about our task by showing a demo video (5 minutes long) of our task interface. During the tool exploration phase, they are invited to use the freely available task interface. We have asked participants to share their screens, which allows us to follow their interaction with the tool. It enables us to provide help when needed. We ensure that all participants achieved a sufficient understanding of the functionality of the tool and support them when we notice fundamental misconceptions about how the tool works and how they can operate it. Each participant has spent approximately 20 minutes in the tool exploration phase. All communications are conducted in English.

For evaluation, we randomly select 150 tweets. Before starting the task, we conduct a pilot test with a separate batch of 10 tweets (excluded from randomly selected 150 tweets). Our pilot test indicates that it takes around < 1 minute to complete a tweet. Each annotator is provided with 50 tweets per batch, completing three batches in total. Participation in the task has been taken place between June and September of 2024.

We calculate the inter-annotator agreement using Krippendorff’s α [42], where $\alpha = 1$ suggests perfect agreement, and $\alpha = 0$ suggests chance-level agreement. We find $\alpha = 0.979$, which suggests a satisfactory level of agreement, indicating a reliable rating.

Method	$Acc_{overall}$	Acc_{MF}	$F1_{MF}$
few-shot _{wo/expl}	64.81%	71.7%	72.58%
few-shot_{w/expl}	90.79%	92.67%	93.51%

Table 2: Morality Frame Identification Task.

4.3 Results

If majority vote from annotators’ choose ‘yes’ on LLMs-generated labels for a text, we consider that a ‘win’ situation. If the majority vote from the annotators’ chooses ‘no’, we calculate the majority vote to get moral foundations and moral role labels provided by the annotators during the human evaluation phase for the corresponding text. We calculate the overall accuracy of the LLMs-generated morality frame prediction task as well as the accuracy and macro average F1 score for the moral foundation prediction task of LLMs. According to human evaluation, few-shot prompting with explanation using LLMs can predict morality frame in 90.79% accuracy. The

Participant	Task Diff. (w/o expl.)	Task Diff. (w/ expl.)	Expl. Helpful?	Red. Cog. Load?	Avg. Time/Batch (min)
P1	very hard	easy	yes	yes	30
P2	very hard	easy	yes	yes	45
P3	very hard	easy	yes	yes	33
P4	hard	easy	yes	yes	70
P5	very hard	easy	yes	yes	42
P6	very hard	easy	yes	yes	31
P7	okay	easy	yes	yes	48
P8	hard	okay	yes	yes	30
P9	hard	easy	yes	yes	90

Table 3: Survey Results on Morality Frame Identification Task. "Task Diff. (w/o expl.)": Task Difficulty without label and explanation, "Task Diff. (w/ expl.)": Task Difficulty with label and explanation, "Expl. Helpful?": Is the explanation helpful, "Red. Cog. Load?": Does explanation reduce cognitive load, "Avg. Time/Batch (min)": Average time per batch.

results are provided in Table 2. In section 5, we provide a discussion about our key findings.

4.4 Survey Results

We conduct a survey to gather annotators’ feedback regarding the effectiveness of explanations and their impact on reducing cognitive load, task difficulty, and the time required for each batch of annotations. Results are provided in Table 3. In subsection 5.2, we provide a discussion regarding participants’ comments.

4.5 Ablation Studies

We provide an ablation study in which annotators provide their judgment regarding morality frames without looking at explanations. We find that the overall accuracy drops at 64.81% (1st row of Table 2) when we don’t provide explanations to the annotators.

4.6 Analysis of Morality Framing

In this section, we dive deeper into the specific ways in which moral foundations and entity roles are framed within the vaccine debate on social media. By analyzing the correlation between different moral foundations, reasons, and stances, we aim to uncover how moral narratives are constructed and how these narratives shape public discourse.

4.6.1 Correlation Analysis of Moral Foundations, Reasons, and Stances.

As there are previously annotated stances (i.e., pro-vax, anti-vax, neutral) for COVID-19 vaccine tweets [57] available, we use it for the analysis of 150 tweets. Pacheco et al. [57] provided the reasons (including the phrases under each reason) why people cite to support or oppose the vaccination debate. In addition, we use other themes named ‘vaccine equity’, ‘vaccine rollout’ borrowed from Islam and Goldwasser [30]. Additional details about the reasons [57] and themes [30] can be found in the original publications. Finally, we utilize those reasons for our analysis. To evaluate the correlation between the different dimensions of the analysis, we calculate the Pearson’s correlation matrices and present them in Fig. 6.

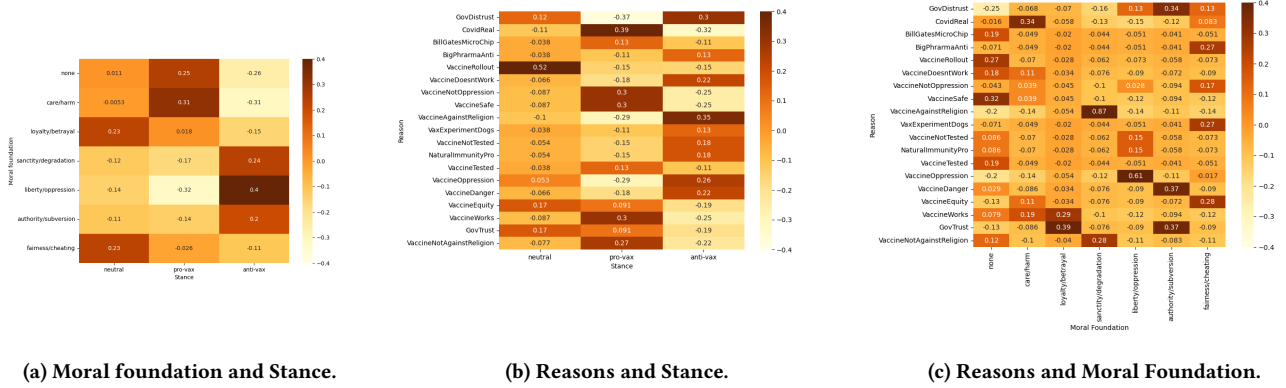


Figure 6: Correlation Heatmaps.

We can interpret reasons as distributions over moral foundations and stances (and vice-versa). This analysis provides a useful way to explain each of these dimensions. For example, we see that **care/harm** is strongly correlated with reasons such as *COVID is real*, *vaccine works* (Fig. 6c). Another expected trend is **liberty/oppression** being highly correlated with *vaccine mandate is oppression* (Fig. 6c). On the other hand, **sanctity/degradation** moral foundation is correlated with *vaccine is against religion*, *vaccine is not against religion* (Fig. 6c). Noticeable correlations are found between **loyalty/betrayal** moral foundation and *trust in government* reason, as well as between **authority/subversion** moral foundation and reasons *vaccine is dangerous*, *trust in government*, and *distrust in government* (Fig. 6c). The **fairness/cheating** moral foundation correlates more with *vaccine equity* theme (Fig. 6c).

In Fig. 6a, we notice that **pro-vaccination** stance correlates with **care/harm** moral foundation and **anti-vaccination** stance correlates with **liberty/oppression** moral foundation. In Fig. 6b, we can see **pro-vaccination** stance strongly correlates with reasons such as *COVID is real*, *vaccine works*, *vaccine is safe*, *vaccine mandate is not oppression*, *vaccine is not against religion*. In contrast, **anti-vaccination** stance highly correlates with reasons such as *vaccine is against religion*, *distrust in government*, *vaccine is oppression*, *vaccine doesn't work*, *vaccine is dangerous* (Fig. 6b).

4.6.2 Moral Foundations and Entity Roles. In addition to examining moral foundations, we analyze how specific entities (e.g., government figures, institutions) are framed within these moral narratives. The roles of entities as actors or targets, and their associated polarity (positive or negative), provide a deeper understanding of how different sides of the debate construct their arguments. In Table 4, we provide the top two reasons for **sanctity/degradation**, **fairness/cheating**, and **authority/subversion** moral foundations. We choose those moral foundations given that they are evenly split among stances and are active for different reasons. We show the top two (**entity, role, polarity**) tuples for each reason. We show that while these moral foundations are used by people on both sides, the reasons offered and entities used vary. On the anti-vaccination side, *authority figures* (i.e., *Fauci*, *Biden*) and *vaccines* are portrayed

MF	Top 2 Reasons for MF	Most Freq. Entity Roles
sanctity/degradation	VaccineAgainstReligion	100% anti-vax (vaccine, actor, negative) (Christian, target, negative)
	VaccineNotAgainstReligion	100% pro-vax (vaccine, actor, positive) (I, target, positive)
fairness/cheating	GovDistrust	67% anti-vax (Fauci, actor, negative) (children, target, negative)
	CovidReal	100% pro-vax (Fox News, actor, negative) (viewers, target, negative)
authority/subversion	GovDistrust	100% anti-vax (Biden, actor, negative) (people, target, negative)
	GovTrust	100% pro-vax (FDA, actor, positive) (Americans, target, positive)

Table 4: Top 2 reasons for sanctity/degradation, fairness/cheating, and authority/subversion moral foundations. Their most frequent opinions and entity roles.

as **negative actors**, while *people*, *children*, and *Christians* are portrayed as **negative targets** (Table 4). On the pro-vaccination side, *Fox News* is portrayed as **negative actor**, and the *viewers* are portrayed as a **negative targets**. Besides, we notice that *FDA*, *vaccine* are portrayed as **positive actors** and *Americans* are portrayed as **positive targets** on the pro-vaccine side (Table 4). This analysis highlights the complexity of moral framing in social media discourse and underscores the importance of context in understanding these narratives. By identifying how different entities are framed within these moral narratives, we can better understand the broader social and cultural dynamics at play.

5 Discussion

This research addresses the question of whether LLMs can effectively assist human annotators in the complex task of identifying

morality frames within social media content, particularly in discussions surrounding vaccination. By structuring our study around two key steps—first, generating necessary data and explanations using LLMs, and second, evaluating these outputs with human input—we aim to understand the potential and limitations of AI.

The task we focus on, moral frame identification, presents unique challenges due to its psycholinguistic nature. By leveraging LLMs' few-shot learning capabilities, augmented with detailed explanations, we explore how AI can reduce the cognitive load on human annotators and improve the efficiency and accuracy of their work. The integration of a "think-aloud" protocol allows us to gather insights into how annotators interact with AI-generated outputs, further enriching our understanding of human-AI collaboration in this context. In the following, we highlight our key findings.

5.1 Key Findings from LLMs' Few-shot Prompting with Explanation Capabilities

Our results in Table 2 suggest that LLMs can identify morality frames using the few-shot prompting with the explanation method with an overall accuracy of **90.79%**. In some cases, annotators mark LLMs prediction *incorrect* though the moral foundation prediction is correct due to the incorrect actor-target role with polarity predictions by LLMs. That's the reason we observe **higher accuracy** (92.67%) and **F1 score** (93.51%) in moral foundation prediction than the overall morality frame prediction. This finding underscores the complexity of moral frame identification, where even advanced AI systems may struggle with nuanced aspects of psycholinguistic tasks.

We find that LLMs can identify **moral cases better than non-moral ('none') cases**. In non-moral cases, we notice that LLMs sometimes show bias in the explanation. For instance, consider the tweet: *"Pentagon to require COVID vaccine for all troops by Sept. 15"*. LLMs classify this under the moral foundation of 'authority/subversion', providing the explanation: *"The text suggests the exercise of authority by the Pentagon (governmental institution) to ensure safety protocols (requiring vaccine) for the troops"*. Here, LLMs depict the 'Pentagon' as a positive actor and 'troops' as a positive target, with the explanation: *"In this situation, the Pentagon (actor) is exercising their authority positively by requiring the COVID-19 vaccine for all troops (target) as a measure to protect their health"*. However, the text itself does not explicitly mention any safety protocols for the troops. Actually, this text is a fact without explicit indication of moral judgment. This indicates that while LLMs can be powerful tools, careful oversight, and human judgment remain essential in tasks that involve moral and ethical considerations.

Morality frame identification is a complex task, and occasionally, we encounter three-way disagreement among annotators. For example, in the following text: *"We were in a better situation this time last year with no vaccine. The gene jab is the equivalent of pouring petrol on the flames, leaky vaccines lead to more variants & the trialists have destroyed their immune systems. Bunch of Turkeys and xmas is coming!"*, LLMs identify moral foundation as 'sanctity/degradation'. One annotator thinks LLMs generated answer is *correct*, and two annotators think it is *incorrect*. Of these dissenting annotators, one identifies the moral foundation of this text as 'none', whereas the other opts for 'care/harm'. We resolve those

issues through discussion. These disagreements highlight the inherent subjectivity in moral frame identification and the need for discussion and consensus-building when using AI tools in qualitative research. Such instances underscore the value of a collaborative framework where AI augments rather than replaces human judgment.

5.2 Key Findings from Participants' Survey

Feedback from the participant survey reinforces the value of integrating LLM-generated labels and explanations into the annotation process. In the feedback survey, five (*P1, P2, P3, P5, P6*) out of nine participants mention the morality frame identification task without having any labels and explanations is **very hard** task. However, they note that the addition of LLMs-generated labels and explanations reduced the task difficulty to **easy** (Table 3). Three participants (*P4, P8, P9*) identify the task itself as a **hard** task. Among them, two (*P4* and *P9*) mention that the provided outputs and explanations reduce the task complexity to **easy**, while participant *P8* rates the task complexity as **okay** (Table 3). On the other hand, participant *P7* believes that the task complexity is **okay** without any labels or explanations, but it is reduced to **easy** with them. Overall, all participants find the morality frame identification task challenging without labels or explanations but significantly easier with the addition of LLMs-generated labels and explanations.

All annotators agree that the labels and explanations provided by LLMs are **extremely helpful** and **significantly reduce their cognitive load**. These explanations are particularly valuable for aiding in the determination of the correct annotations for identifying actor-target roles with polarity. Annotators find that the explanations from LLMs are accurate in most instances. Even when inaccuracies occurred, the explanations still facilitate correct adjustments to the annotations. This finding is crucial for the human-computer interaction (HCI) community, as it highlights how AI tools can alleviate the mental burden associated with complex tasks. The accuracy of these explanations in guiding correct annotations also points to the potential of LLMs to support rather than supplant human expertise in qualitative research.

Significantly, the presence of explanations **reduces the time** required to complete the same annotations **by at least 50%**. This efficiency not only boosts annotators' confidence in the accuracy of their work but also alleviates the confusion that often arose in tasks lacking explanations. The availability of initial answers through these explanations prove immensely beneficial in streamlining the annotation process. Furthermore, annotators note a **significant decrease in exhaustion—over 60% less**—after processing batches with explanations compared to those without, underscoring the substantial cognitive relief provided by the initial guidance from LLMs. This finding underscores the value of AI in enhancing both the quality and sustainability of human annotation efforts.

All participants recognize our tool as a significant enhancement to their research practices. Annotators perceive these computational tools as facilitators of collaboration rather than full automation, reflecting the growing body of research on human-AI collaboration in qualitative research contexts [5, 10, 14, 21, 36].

6 Conclusion and Future Work

Our study has validated the significant role that LLMs can play in supporting annotators during complex psycholinguistic tasks such as morality frame identification within the context of vaccination debates on social media. We prompt LLMs with few-shot examples and explanations to identify moral foundation and actor-target-polarity with corresponding explanations. By integrating LLMs-generated answers with explanations into the annotation process, our study shows whether and to what extent it will aid annotators and streamline workflow. Our research offers insights into how LLMs can assist annotators and lays the groundwork for future research in generating data and explanations for complex psycholinguistic tasks.

However, the study also underscores the importance of human oversight in AI-assisted tasks, particularly in areas requiring nuanced understanding and interpretation. The occasional misclassifications and biases introduced by the LLMs highlight the need for a collaborative approach where AI serves to augment rather than replace human judgment.

Moving forward, this research opens up several avenues for further exploration. In the future, we aim to add corresponding themes and arguments [31, 33, 58, 59] related to the vaccine debate while prompting LLMs to further improve the accuracy of the morality frame identification task [30]. Though we show the effectiveness of our framework focusing on the vaccine debate on social media, it is a domain-agnostic framework. Future studies could test the effectiveness of LLMs across various domains, such as political discourse, climate debate etc. Besides, future research should investigate on the sustainable adoption of LLMs in research settings to aid annotators. There is a need to investigate the biases that LLMs may introduce further.

In conclusion, our work contributes to the growing body of research on human-AI collaboration. It demonstrates the potential of LLMs to support and enhance human decision-making in complex annotation tasks, offering a promising direction for future research in AI-assisted qualitative analysis. As AI continues to evolve, its role in augmenting human capabilities will likely expand, making it an increasingly integral part of the HCI landscape.

7 Limitations

LLMs are pre-trained on a huge amount of human-generated text. As a result, they may inherently contain many human biases [8, 11]. We did not consider any bias in our task.

A previous study by Johnson and Goldwasser [37] has shown that a single tweet may contain multiple moral foundations. We did not consider multi-label moral foundations in this work.

We are aware of the limitations due to potential biases arising from our study participant selection. Our study relied on voluntary participation, which might induce bias as potential participants would agree to participate if they had at least some degree of interest and openness toward the psycholinguistic task. Additionally, in technology-based studies, “novelty bias” could influence the study outcome when participants are excited to try something new [49]. Hence, we made sure that our sample included participants with different age groups, academic levels, diverse research areas, and gender (see subsection 4.2.1).

The main objective of this paper is to determine whether labels and explanations generated by LLMs are helpful and can assist annotators on hard psycholinguistic tasks. However, the approach is designed to be scalable without any modifications.

Acknowledgments

We are thankful for the reviews of anonymous reviewers. This work was partially funded by the Purdue Graduate School Summer Research Grant (to TI) and the NSF CAREER award IIS-2048001.

References

- [1] Woo-kyoung Ahn, William F Brewer, and Raymond J Mooney. 1992. Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory, and Cognition* (1992).
- [2] Lola Al-Uqdah, F Abbron Franklin, Chu-Chuan Chiu, and Brianna N Boyd. 2022. Associations between social media engagement and vaccine hesitancy. *Journal of Community Health* 47, 4 (2022), 577–587.
- [3] Avnika B Amin, Robert A Bednarczyk, Cara E Ray, Kala J Melchiori, Jesse Graham, Jeffrey R Huntsinger, and Saad B Omer. 2017. Association of moral values with vaccine hesitancy. *Nature Human Behaviour* 1, 12 (2017), 873–880.
- [4] Pablo Barberá, Anita R Gohdes, Evgeniia Iakhnis, and Thomas Zeitzoff. 2024. Distract and divert: How world leaders use social media during contentious politics. *The International Journal of Press/Politics* 29, 1 (2024), 47–73.
- [5] Eric PS Baumer et al. 2020. Topicalizer: reframing core concepts in machine learning visualization by co-designing for interpretivist scholarship. *Human-Computer Interaction* 35, 5-6 (2020), 452–480.
- [6] Mariano Gastón Beiró, Jacopo D’Ignazi, Victoria Perez Bustos, María Florencia Prado, and Kyriaki Kalimeri. 2023. Moral narratives around the vaccination debate on facebook. In *Proceedings of the ACM Web Conference 2023*. 4134–4141.
- [7] Sandeep Bhanot. 2012. Use of social media by companies to reach their customers. *SIES Journal of Management* 8, 1 (2012).
- [8] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (Technology) is Power: A Critical Survey of “Bias” in NLP. In *ACL*.
- [9] Erika Bonnevise, Allison Gallegos-Jeffrey, Jaclyn Goldberg, Brian Byrd, and Joseph Smyser. 2021. Quantifying the rise of vaccine opposition on Twitter during the COVID-19 pandemic. *Journal of communication in healthcare* 14, 1 (2021), 12–19.
- [10] Petter Bae Brandtzaeg, Marita Skjuve, and Asbjørn Følstad. 2022. My AI friend: How users of a social chatbot understand their human-AI friendship. *Human Communication Research* 48, 3 (2022), 404–429.
- [11] Tom Brown et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* (2020).
- [12] Joel Brynielsson et al. 2018. Informing crisis alerts using social media: Best practices and proof of concept. *Journal of contingencies and crisis management* (2018).
- [13] Eugene Y Chan. 2021. Moral foundations underlying behavioral compliance during the COVID-19 pandemic. *Personality and individual differences* (2021).
- [14] Nan-Chen Chen et al. 2018. Using machine learning to support qualitative coding in social science: Shifting the focus to ambiguity. *ACM TIS* (2018).
- [15] Wen-ying Sylvia Chou, Yvonne M Hunt, Ellen B Beckjord, Richard P Moser, and Bradford W Hesse. 2009. Social media use in the United States: implications for health communication. *Journal of medical Internet research* 11, 4 (2009), e1249.
- [16] Aakanksha Chowdhery, Sharan Narang, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* (2023).
- [17] Rodrigo Diaz and Florian Cova. 2021. Reactance, morality, and disgust: The relationship between affective dispositions and compliance with official health recommendations during the COVID-19 pandemic. *Cognition and Emotion* (2021).
- [18] Bosheng Ding, Chengwei Qin, Linlin Liu, Yew Ken Chia, Boyang Li, Shafiq Joty, and Lidong Bing. 2023. Is GPT-3 a Good Data Annotator?. In *ACL*.
- [19] Emilio Ferrara, Herbert Chang, Emily Chen, Goran Muric, and Jaimin Patel. 2020. Characterizing social media manipulation in the 2020 US presidential election. *First Monday* (2020).
- [20] Emilio Ferrara, Stefano Cresci, and Luca Luceri. 2020. Misinformation, manipulation, and abuse on social media in the era of COVID-19. *Journal of Computational Social Science* 3 (2020), 271–277.
- [21] Simret Araya Gebreegziabher et al. 2023. Patat: Human-ai collaborative qualitative coding with explainable interactive rule synthesis. In *CHI*.
- [22] Rebekah Getman, Mohammad Helmi, Hal Roberts, Alfa Yansane, David Cutler, and Brittany Seymour. 2018. Vaccine hesitancy and online information: the influence of digital networks. *Health Education & Behavior* 45, 4 (2018), 599–606.
- [23] Fabrizio Gilardi, Meysam Alizadeh, and Maël Kubli. 2023. ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences* 120, 30 (2023), e2305016120.

- [24] Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *JPS* (2009).
- [25] Jonathan Haidt and Jesse Graham. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research* 20, 1 (2007), 98–116.
- [26] Jonathan Haidt and Craig Joseph. 2004. Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus* 133, 4 (2004), 55–66.
- [27] Raquel G Hernandez et al. 2021. The COVID-19 vaccine social media infodemic: healthcare providers' missed dose in addressing misinformation and vaccine hesitancy. *Human Vaccines & Immunotherapeutics* (2021).
- [28] Gabriela Hoefler et al. 2022. Bridging the Social Distance: Offline to Online Social Support during the COVID-19 Pandemic. *CSCW* (2022).
- [29] Fan Huang, Haewoon Kwak, and Jisun An. 2023. Is chatgpt better than human annotators? potential and limitations of chatgpt in explaining implicit hate speech. In *Companion proceedings of the ACM web conference 2023*. 294–297.
- [30] Tunazzina Islam and Dan Goldwasser. 2022. Understanding COVID-19 Vaccine Campaign on Facebook using Minimal Supervision. In *2022 IEEE International Conference on Big Data (Big Data)*. IEEE, 585–595.
- [31] Tunazzina Islam and Dan Goldwasser. 2024. Discovering Latent Themes in Social Media Messaging: A Machine-in-the-Loop Approach Integrating LLMs. *arXiv preprint arXiv:2403.10707* (2024).
- [32] Tunazzina Islam and Dan Goldwasser. 2024. Post-hoc Study of Climate Microtargeting on Social Media Ads with LLMs: Thematic Insights and Fairness Evaluation. *arXiv preprint arXiv:2410.05401* (2024).
- [33] Tunazzina Islam and Dan Goldwasser. 2024. Uncovering Latent Arguments in Social Media Messaging by Employing LLMs-in-the-Loop Strategy. *arXiv preprint arXiv:2404.10259* (2024).
- [34] Tunazzina Islam, Shamik Roy, and Dan Goldwasser. 2023. Weakly Supervised Learning for Analyzing Political Campaigns on Facebook. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 17. 411–422.
- [35] Tunazzina Islam, Ruqi Zhang, and Dan Goldwasser. 2023. Analysis of Climate Campaigns on Social Media Using Bayesian Model Averaging. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*.
- [36] Jialun Aaron Jiang, Kandrea Wade, Casey Fiesler, and Jed R Brubaker. 2021. Supporting serendipity: Opportunities and challenges for Human-AI Collaboration in qualitative analysis. *CSCW* (2021).
- [37] Kristen Johnson and Dan Goldwasser. 2018. Classification of moral foundations in microblog political discourse. In *ACL*.
- [38] S Mo Jones-Jang and Myojung Chung. 2024. Can we blame social media for polarization? Counter-evidence against filter bubble claims during the COVID-19 pandemic. *New Media & Society* 26, 6 (2024), 3370–3389.
- [39] Minsol Kim, Aliea L Nallbani, and Abby Rayne Stovall. 2024. Exploring LLM-based Chatbot for Language Learning and Cultivation of Growth Mindset. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*.
- [40] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems* 35 (2022), 22199–22213.
- [41] Mahi Kolla, Siddharth Salunkhe, Eshwar Chandrasekharan, and Koustuv Saha. 2024. LLM-Mod: Can Large Language Models Assist Content Moderation?. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*.
- [42] Klaus Krippendorff. 2004. Measuring the reliability of qualitative text analysis data. *Quality and quantity* 38 (2004), 787–800.
- [43] Andrew Lampinen et al. 2022. Can language models learn from explanations in context?. In *EMNLP Findings*.
- [44] Teven Le Scao, Angela Fan, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. (2022).
- [45] Susan Lin, Jeremy Warner, et al. 2024. Rambler: Supporting Writing With Speech via LLM-Assisted Gist Manipulation. In *CHI*.
- [46] Tania Lombrozo. 2006. The structure and function of explanations. *Trends in cognitive sciences* 10, 10 (2006), 464–470.
- [47] Tania Lombrozo and Susan Carey. 2006. Functional explanation and the function of explanation. *Cognition* (2006).
- [48] Xinhao Mei et al. 2024. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (2024).
- [49] Joy Ming, Sharon Heung, Shiri Azenkot, and Aditya Vashistha. 2021. Accept or address? Researchers' perspectives on response bias in accessibility research. In *ACM SIGACCESS Conference on Computers and Accessibility*.
- [50] Swaroop Mishra et al. 2022. Reframing Instructional Prompts to GPTk's Language. In *ACL Findings*.
- [51] Swaroop Mishra, Daniel Khashabi, Chitta Baral, and Hannaneh Hajishirzi. 2022. Cross-Task Generalization via Natural Language Crowdsourcing Instructions. In *ACL*.
- [52] Patricia Moravec, Randall Minas, and Alan R Dennis. 2018. Fake news on social media: People believe what they want to believe when it makes no sense at all. *Kelley School of Business research paper* 18-87 (2018).
- [53] Marcia Mundt, Karen Ross, and Charla M Burnett. 2018. Scaling social movements through social media: The case of Black Lives Matter. *Social media+ society* 4, 4 (2018), 2056305118807911.
- [54] OpenAI. 2024. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024. (2024).
- [55] Daniela Orr, Ayelet Baram-Tsabari, and Keren Landsman. 2016. Social media as a platform for health-related public debates and discussions: the Polio vaccine on Facebook. *Israel journal of health policy research* 5 (2016), 1–11.
- [56] Maria Leonor Pacheco and Dan Goldwasser. 2021. Modeling Content and Context with Deep Relational Learning. *TACL* (2021).
- [57] Maria Leonor Pacheco, Tunazzina Islam, Monal Mahajan, Andrey Shor, Ming Yin, Lyle Ungar, and Dan Goldwasser. 2022. A Holistic Framework for Analyzing the COVID-19 Vaccine Debate. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 5821–5839.
- [58] Maria Leonor Pacheco, Tunazzina Islam, Lyle Ungar, Ming Yin, and Dan Goldwasser. 2022. Interactively uncovering latent arguments in social media platforms: A case study on the covid-19 vaccine debate. In *Fourth Workshop on Data Science with Human-in-the-Loop*.
- [59] Maria Leonor Pacheco, Tunazzina Islam, Lyle Ungar, Ming Yin, and Dan Goldwasser. 2023. Interactive Concept Learning for Uncovering Latent Themes in Large Text Collections. In *Findings of the Association for Computational Linguistics: ACL 2023*.
- [60] Stefano Pagliaro et al. 2021. Trust predicts COVID-19 prescribed and discretionary behavioral intentions in 23 countries. *PLoS one* (2021).
- [61] Neha Puri, Eric A Coomes, Hourmazd Haghighyan, and Keith Gunaratne. 2020. Social media and vaccine hesitancy: new updates for the era of COVID-19 and globalized infectious diseases. *Human vaccines & immunotherapeutics* (2020).
- [62] Yasmim Mendes Rocha et al. 2021. The impact of fake news on social media and its influence on health during the COVID-19 pandemic: A systematic review. *Journal of Public Health* (2021), 1–10.
- [63] Jon Roozenbeek, Claudia R Schneider, et al. 2020. Susceptibility to misinformation about COVID-19 around the world. *R. Soc. Open Sci.* (2020).
- [64] Isabel Rossen, Mark J Hurlstone, Patrick D Dunlop, and Carmen Lawrence. 2019. Accepters, fence sitters, or rejecters: Moral profiles of vaccination attitudes. *Social Science & Medicine* 224 (2019), 23–27.
- [65] Shamik Roy, Nishanth Sridhar Nakshatri, and Dan Goldwasser. 2022. Towards Few-Shot Identification of Morality Frames using In-Context Learning. In *Proceedings of the Fifth Workshop on NLP+ CSS*.
- [66] Shamik Roy, Maria Leonor Pacheco, and Dan Goldwasser. 2021. Identifying Morality Frames in Political Tweets using Relational Learning. In *EMNLP*.
- [67] Arnaldo Santoro et al. 2023. Analyzing the changing landscape of the Covid-19 vaccine debate on Twitter. *Social Network Analysis and Mining* (2023).
- [68] Ana Lucia Schmidt, Fabiana Zollo, Antonio Scala, Cornelia Betsch, and Walter Quattrocchi. 2018. Polarization of the vaccination debate on Facebook. *Vaccine* 36, 25 (2018), 3606–3612.
- [69] Eva Sharma, Koustuv Saha, Sindhu Kiranmai Erala, Sucheta Ghoshal, and Munmun De Choudhury. 2017. Analyzing ideological discourse on social media: A case study of the abortion debate. In *International conference of the computational social science society of the Americas*.
- [70] Laura M Smith, Linhong Zhu, Kristina Lerman, and Zornitsa Kozareva. 2013. The role of social media in the discussion of controversial topics. In *2013 International Conference on Social Computing*. IEEE, 236–243.
- [71] Fabio Tagliabue, Luca Galassi, and Pierpaolo Mariani. 2020. The “pandemic” of disinformation in COVID-19. *SN comprehensive clinical medicine* (2020).
- [72] Kathie M d'I Trean, Hywel TP Williams, and Saffron J O'Neill. 2020. Online misinformation about climate change. *Wiley Interdisciplinary Reviews: Climate Change* 11, 5 (2020), e665.
- [73] Shuohang Wang, Yang Liu, Yichong Xu, Chenguang Zhu, and Michael Zeng. 2021. Want To Reduce Labeling Cost? GPT-3 Can Help. In *EMNLP Findings*.
- [74] Yuping Wang, Chen Ling, and Gianluca Stringhini. 2023. Understanding the use of images to spread COVID-19 misinformation on Twitter. *CSCW* (2023).
- [75] Jason Wei et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* (2022).
- [76] Maxwell A Weinzierl and Sanda M Harabagiu. 2022. From hesitancy framings to vaccine hesitancy profiles: A journey of stance, ontological commitments and moral foundations. In *ICWSM*.
- [77] Liang Wu, Fred Morstatter, Kathleen M Carley, and Huan Liu. 2019. Misinformation in social media: definition, manipulation, and detection. *ACM SIGKDD explorations newsletter* (2019).
- [78] Xiaoteng Xu, Jiayu Yin, et al. 2024. Jamplate: Exploring LLM-Enhanced Templates for Idea Reflection. In *IUI*.
- [79] Zhiying Yue, Renwen Zhang, and Jun Xiao. 2022. Passive social media use and psychological well-being during the COVID-19 pandemic: The role of social comparison and emotion regulation. *Computers in Human Behavior* (2022).
- [80] Renwen Zhang, Natalya N. Bazarova, and Madhu Reddy. 2021. Distress disclosure across social media platforms during the COVID-19 pandemic: Untangling the effects of platforms, affordances, and audiences. In *CHI*.
- [81] Wei Zhao et al. 2023. Older adults using technology for meaningful activities during COVID-19: An analysis through the lens of self-determination theory. In *CHI*.