

# Community-Based Fact-Checking on Twitter’s Birdwatch Platform

Nicolas Pröllochs

JLU Giessen  
 nicolas.proellocks@wi.jlug.de

## Abstract

Misinformation undermines the credibility of social media and poses significant threats to modern societies. As a countermeasure, Twitter has recently introduced “Birdwatch,” a community-driven approach to address misinformation on Twitter. On Birdwatch, users can identify tweets they believe are misleading, write notes that provide context to the tweet and rate the quality of other users’ notes. In this work, we empirically analyze how users interact with this new feature. For this purpose, we collect all Birdwatch notes and ratings between the introduction of the feature in early 2021 and end of July 2021. We then map each Birdwatch note to the fact-checked tweet using Twitter’s historical API. In addition, we use text mining methods to extract content characteristics from the text explanations in the Birdwatch notes (e. g., sentiment). Our empirical analysis yields the following main findings: (i) users more frequently file Birdwatch notes for misleading than not misleading tweets. These misleading tweets are primarily reported because of factual errors, lack of important context, or because they treat unverified claims as facts. (ii) Birdwatch notes are more helpful to other users if they link to trustworthy sources and if they embed a more positive sentiment. (iii) The social influence of the author of the source tweet is associated with differences in the level of user consensus. For influential users with many followers, Birdwatch notes yield a lower level of consensus among users and community-created fact checks are more likely to be seen as being incorrect and argumentative. Altogether, our findings can help social media platforms to formulate guidelines for users on how to write more helpful fact checks. At the same time, our analysis suggests that community-based fact-checking faces challenges regarding opinion speculation and polarization among the user base.

## Introduction

Misinformation on social media has become a major focus of public debate and academic research. Previous works have demonstrated that misinformation is widespread on social media platforms and that misinformation diffuses significantly farther, faster, deeper, and more broadly than the truth (e.g., Vosoughi, Roy, and Aral; Pröllochs, Bär, and Feuerriegel; Pröllochs, Bär, and Feuerriegel 2018; 2021b; 2021a). Concerns about misinformation on social media

have been rising in recent years, particularly given its potential impacts on elections (Aral and Eckles 2019; Bakshy, Messing, and Adamic 2015), public health (Solovev and Pröllochs 2022b), and public safety (Starbird 2017). Major social media providers (e. g., Twitter, Facebook) thus have been called upon to develop effective countermeasures to combat the spread of misinformation on their platforms (Lazer et al. 2018; Pennycook et al. 2021).

A crucial prerequisite to curb the spread of misinformation on social media is its accurate identification (Pennycook and Rand 2019a). Predominant approaches to identify misinformation on social media can be grouped into two categories. First, human-based systems in which (human) experts or fact-checking organizations (e. g., snopes.com, politifact.com, factcheck.org) determine the veracity (Hasan et al. 2017; Shao et al. 2016). Second, machine learning-based systems can automatically classify veracity (Ma et al. 2016; Qazvinian et al. 2011). Here machine learning models are typically trained to classify misinformation using content-based features (e. g. text, images, video), context-based features (e. g., time, location), or based on propagation patterns (i. e., how misinformation circulates among users). Yet both approaches have inherent drawbacks: (i) expert’s verification tends to be accurate but is difficult to scale given the limited number of professional fact-checkers (Pennycook and Rand 2019a). (ii) Machine learning-based detection is scalable but the prediction performance tends to be unsatisfactory (Wu et al. 2019). Complementary approaches are thus necessary to identify misinformation on social media both accurately and at scale.

As an alternative, recent research has proposed to build on collective intelligence and the “wisdom of crowds” to fact-check social media content (Micallef et al. 2020; Bhuiyan et al. 2020; Pennycook and Rand 2019a; Epstein, Pennycook, and Rand 2020; Allen et al. 2020, 2021; Godel et al. 2021). The wisdom of crowds is the phenomenon that when many individuals independently make an assessment, the aggregate user judgments will be closer to the truth than most individual estimates or even experts (Frey and van de Rijdt 2021). Applying the concept of crowd wisdom to fact-checking is appealing as it would allow for large numbers of fact-checks that can be inexpensively and frequently acquired (Allen et al. 2021; Pennycook and Rand 2019a). However, research and earlier attempts to harness the wis-

dom of crowds to identify misleading social media content have so far produced mixed results. On the one hand, experimental evidence suggest that the assessment of even small crowds is comparable to those from experts (Allen et al. 2021; Bhuiyan et al. 2020; Pennycook and Rand 2019a). On the other hand, existing crowd-based fact-checking initiatives such as *TruthSquad*, *Factcheck.EU*, and *WikiTribune* had limited success and did not prove to be a model to produce high-quality fact-checks at scale (Bhuiyan et al. 2020). Due to quality issues with crowd-created fact-checks, these initiatives required workarounds involving the assessment of experts – which again limited their scalability (Bhuiyan et al. 2020; Bakabar 2018). In sum, it has been found to be challenging to implement real-world community-based fact-checking platforms that uphold high quality and scalability.

Informed by these findings, Twitter has recently launched “Birdwatch,” a new attempt to address misinformation on social media by harnessing the wisdom of crowds. Different from earlier crowd-based fact-checking initiatives, Birdwatch is a community-driven approach to identify misleading tweets directly on Twitter (see example in Fig. 1). The idea is that the user base on Twitter provides a wealth of knowledge that can help in the fight against misinformation. On Birdwatch, users can identify tweets they believe are misleading (e. g., factual errors) or not misleading (e. g., satire) and write (textual) notes that provide *context* to the tweet. They can also add links to their sources of information. A key feature of Birdwatch is that it implements a rating mechanism that allows users to rate the quality of other participants’ notes. These ratings should help to identify the context that people will find most helpful and raise its visibility to other users. While the Birdwatch feature is currently in pilot phase and only directly visible on tweets to pilot participants in the U. S., Twitter’s goal is that Birdwatch will be available to everyone on Twitter.

**Research goal:** In this work, we provide a holistic analysis of the Birdwatch pilot on Twitter. We empirically analyze how users interact with Birdwatch and study factors that make community-created fact checks more likely to be perceived as helpful or unhelpful by other users. Specifically, we address the following research questions:

- (RQ1) *What are specific reasons due to which Birdwatch users report tweets?*
- (RQ2) *How do Birdwatch notes for tweets categorized as being misleading vs. not misleading differ in terms of their content characteristics (e. g., sentiment, length)?*
- (RQ3) *Are tweets from Twitter accounts with certain characteristics (e. g., politicians, accounts with many followers) more likely to be fact-checked on Birdwatch?*
- (RQ4) *Which characteristics of Birdwatch notes are associated with greater helpfulness for other users?*
- (RQ5) *How is the level of consensus among users associated with the social influence of the author of the source tweet?*

**Methodology:** To address our research questions, we collect *all* Birdwatch notes and ratings from the Birdwatch website between the introduction of the feature on January

23, 2021, and the end of July 2021. This comprehensive dataset contains 11,802 Birdwatch notes and 52,981 ratings. We use text mining methods to extract content characteristics from the text explanations of the Birdwatch notes (e. g. sentiment, length). In addition, we employ the Twitter historical API to collect information about the source tweets (i. e., the fact-checked tweets) referenced in the Birdwatch notes. We then perform an empirical analysis of the observational data and use regression models to understand how (i) the user categorization, (ii) the content characteristics of the text explanation, and (iii) the social influence of the source tweet are associated with the helpfulness of Birdwatch notes.

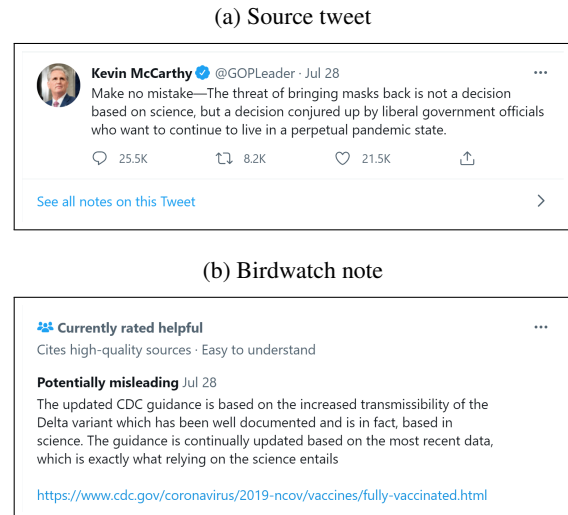


Figure 1: Screenshot of an exemplary community-created fact-check on Birdwatch.

**Contributions:** To the best of our knowledge, this study is the first to present a thorough empirical analysis of Twitter’s Birdwatch feature. In contrast to earlier studies focusing on whether crowds are *able* to accurately assess social media content, we contribute to research into misinformation and crowdsourced fact-checking by shedding light on how users *interact* with a community-based fact-checking system. Our work yields the following main findings:

1. Users file a larger number of Birdwatch notes for misleading than for not misleading tweets. Misleading tweets are primarily reported because of factual errors, lack of important context or because they treat unverified claims as fact. Not misleading tweets are primarily reported because they are perceived as factually correct or because they clearly refer to personal opinion or satire.
2. Birdwatch notes filed for misleading vs. not misleading tweets differ in terms of their content characteristics. For tweets reported as misleading, authors of Birdwatch notes use a more negative sentiment, more complex language, and tend to write longer text explanations.
3. Tweets from influential politicians (on both sides of the political spectrum in the U. S.) are more likely to be fact-

checked on Birdwatch. For many of these accounts, Birdwatch users overwhelmingly report misinformation.

4. Birdwatch notes reporting misleading tweets are perceived as being more helpful by other Birdwatch users. Birdwatch notes are perceived as being particularly helpful if they provide trustworthy sources and embed a positive sentiment.
5. The social influence of the author of the source tweet is associated with differences in the level of user consensus. For influential users with many followers, Birdwatch notes receive more total votes (helpful & unhelpful) but a lower helpfulness ratio (i. e., a lower level of consensus). Also, Birdwatch notes for influential users are particularly likely to be seen as incorrect and argumentative.

**Implications:** Our findings have direct implications for the Birdwatch platform and future attempts to implement community-based approaches to combat misinformation on social media. We show that users perceive a relatively high share of community-created fact checks as being informative, clear and helpful. Here we find that encouraging users to provide context (e. g., by linking trustworthy sources) and avoid the use of inflammatory language is crucial for users to perceive community-based fact checks as helpful. Despite showing promising potential, our analysis suggests that Birdwatch’s community-driven approach faces challenges concerning opinion speculation and polarization among the user base – in particular with regards to influential accounts.

## Background

### Misinformation on Social Media

Social media has become a prevalent platform for consuming and sharing information online (Bakshy, Messing, and Adamic 2015). It is estimated that almost 62% of the adult population consume news via social media, and this proportion is expected to increase further (Pew Research Center 2016). As any user can share information, quality control for the content has essentially moved from trained journalists to regular users (Kim and Dennis 2019). The inevitable lack of oversight from experts makes social media vulnerable to the spread of misinformation (Shao et al. 2016). Social media platforms have indeed been observed to be a medium that disseminates vast amounts of misinformation (Vosoughi, Roy, and Aral 2018). Several works have studied diffusion characteristics of misinformation on social media (Friggeri et al. 2014; Vosoughi, Roy, and Aral 2018), finding that misinformation spreads significantly farther, faster, deeper, and more broadly than the truth. The presence of misinformation on social media has detrimental consequences on how opinions are formed and the offline world (Bakshy, Messing, and Adamic 2015; Del Vicario et al. 2016). Misinformation thus threatens not only the reputation of individuals and organizations, but also society at large.

Previous research has identified several reasons why misinformation is widespread on social media. On the one hand, misinformation pieces are often intentionally written to mislead other users (Wu et al. 2019). It is thus difficult for users

to spot misinformation from the content itself. On the other hand, the vast majority of social media users do not fact-check articles they read (Geeng, Yee, and Roesner 2020; Vo and Lee 2018). This indicates that social media users are often in a hedonic mindset and avoid cognitive reasoning such as verification behavior (Moravec, Minas, and Dennis 2019). A recent study further suggests that the current design of social media platforms may discourage users from reflecting on accuracy (Pennycook et al. 2021).

The spread of misinformation can also be seen as a social phenomenon. Online social networks are characterized by homophily (McPherson, Smith-Lovin, and Cook 2001), (political) polarization (Levy 2021), and echo chambers (Barberá et al. 2015). In these information environments with low content diversity and strong social reinforcement, users tend to selectively consume information that shares similar views or ideologies while disregarding contradictory arguments (Ecker, Lewandowsky, and Tang 2010). These effects can even be exaggerated in the presence of repeated exposure: once misinformation has been absorbed, users are less likely to change their beliefs even when the misinformation is debunked (Pennycook, Cannon, and Rand 2018).

### Identification of Misinformation

**Identification via expert’s verification:** Misinformation can be manually detected by human experts. In traditional media (e. g., newspapers), this task is typically carried out by professional journalists, editors, and fact-checkers who verify information before it is published (Kim and Dennis 2019). With the possibility for everyone to share information, social media essentially transfers quality control to platform users (Kim and Dennis 2019). This has given rise to a number of third-party fact-checking organizations such as snopes.com and politifact.com that thoroughly investigate and debunk rumors (Wu et al. 2019; Vosoughi, Roy, and Aral 2018). The fact-checking assessments are consumed and broadcast by social media users like any other type of news content and are supposed to help users to identify misinformation on social media (Shao et al. 2016). However, the experts’ verification approach has inherent drawbacks: (i) it does not scale to the volume and speed of content that is generated in social media. Users can create misleading content at a much faster rate than fact-checkers can evaluate it. Misinformation is thus likely to go unnoticed during its period of peak virality (Epstein, Pennycook, and Rand 2020). (ii) Approximately 50% of Americans overall think that fact-checkers are biased and distrust fact-checking corrections (Poynter 2019). Hence even when users become aware of the fact-checks of fact-checking organizations, their impact may be limited by lack of trust. In fact, Vosoughi et al. 2018 show that *fact-checked* false rumors are more viral than the truth.

**Identification using machine learning methods:** Previous research has intensively focused on the problem of misinformation detection via means of supervised machine learning (Ducci, Kraus, and Feuerriegel 2020; Vosoughi, Mohsenvand, and Roy 2017). Researchers typically collect posts and their labels from social media platforms, and then train a classifier based on diverse sets of features. For in-

stance, content-based approaches aim at directly detecting misinformation based on its content, such as text, images, and video. However, different from traditional text categorization tasks, misinformation posts are deliberately made seemingly real and accurate. The predictive power of the actual content in detecting misinformation is thus limited (Wu et al. 2019). Other common information sources include context-based features (e. g., time, location), or propagation patterns (i. e., how misinformation circulates among users) (e.g., Kwon and Cha 2014). Yet also with these features, the lack of ground truth labels poses serious challenges when training a machine learning classifier – in particular in the early stages of the spreading process (Wu et al. 2019).

**Identification through wisdom of crowds:** Recent research has proposed to outsource fact-checking of misinformation to non-expert fact-checkers in the crowd (Micallef et al. 2020; Bhuiyan et al. 2020; Pennycook and Rand 2019a; Epstein, Pennycook, and Rand 2020; Allen et al. 2020, 2021). The idea is to identify misleading social media content by harnessing the wisdom of crowds (Woolley et al. 2010). The wisdom of crowds has been repeatedly observed in a wide range of settings, including online platforms such as Wikipedia and Stack Overflow, where the crowd ensures relatively trustworthy and high-quality accumulation of knowledge (Okoli et al. 2014). Applying the concept of crowd wisdom to fact-checking on social media may have crucial advantages: (i) compared to the expert’s verification approach, which is limited by the number of professional fact-checkers, crowd-based approaches allow to identify misinformation at a large scale (Pennycook and Rand 2019a). (ii) Community-based fact-checking addresses the problem that many users distrust professional fact-checks (Poynter 2019). (iii) The wisdom of crowds literature suggests that, even if the ratings of individual users are noisy and ineffective, in the aggregate user judgments may be highly accurate (Woolley et al. 2010). Experimental studies found that the crowd can in fact be quite accurate in identifying misleading social media content. Here the assessment of even relatively small crowds is comparable to those from experts (Bhuiyan et al. 2020; Epstein, Pennycook, and Rand 2020; Pennycook and Rand 2019a).

Yet, while the crowd may be generally *able* to accurately identify misinformation, not all users may always *choose* to do so (Epstein, Pennycook, and Rand 2020). Important challenges include manipulation attempts (Luca and Zervas 2016), lack of engagement in cognitive reasoning (Pennycook and Rand 2019b), and politically motivated reasoning (Kahan 2017). Each of these factors may hinder an effective fact-checking system. For example, users could purposely try to game the fact-checking system by reporting social media posts as being misleading that do not align with their ideology (irrespective of the perceived veracity) or to achieve partisan ends (Luca and Zervas 2016). Furthermore, the high level of (political) polarization of social media users (Conover et al. 2011; Barberá et al. 2015; Solovey and Pröllochs 2022a), can result in vastly different interpretations of facts or even entirely different sets of acknowledged facts (Ojala et al. 2021). Earlier crowd-based fact-checking initiatives such as *TruthSquad*, *Factcheck.EU*, and

*WikiTribune* (O’Riordan et al. 2019; Florin 2010) indeed experienced significant quality-issues with community-based fact-checks (Bhuiyan et al. 2020; Bakabar 2018). These initiatives required workarounds or hybrid approaches involving final judgments by experts or delegating primary research to experts and secondary tasks to the crowd (Bhuiyan et al. 2020; Bakabar 2018). This suggests that it is challenging to implement real-world community-based fact-checking platforms that uphold high quality and scalability. More precisely, as pointed out by Epstein, Pennycook, and Rand (2020), the challenges that must be overcome in order to implement it successfully are social in nature rather than technical, and thus involve empirical questions about how people *interact* with community-based fact-checking systems.

## What is Birdwatch?

On January 23, 2021, Twitter has launched the “Birdwatch” feature, a new approach to address misinformation on social media by harnessing the wisdom of crowds. In contrast to earlier crowd-based initiatives to fact-checking, Birdwatch is a community-driven approach to identify misleading tweets directly on Twitter. On Birdwatch, users can identify tweets they believe are misleading and write (textual) notes that provide *context* to the tweet. Birdwatch also features a rating mechanism that allows users to rate the quality of other users’ notes. These ratings are supposed to help to identify the context that people will find most helpful and raise its visibility to other users. The Birdwatch feature is currently in the pilot phase in the U. S. and only pilot participants can see community-written fact-checks directly on tweets when browsing Twitter. Users not participating in the pilot can access Birdwatch via a separate Birdwatch website<sup>1</sup>. Twitter’s goal is that community-written notes will be visible directly on tweets, available to everyone on Twitter.

**Birdwatch notes:** Users can add Birdwatch notes to any tweet they come across and think might be misleading. Notes are composed of (i) multiple-choice questions that allow users to state why a tweet might or might not be misleading; (ii) an open text field where users can explain their judgment, as well as link to relevant sources. The maximum number of characters in the text field is 280. Here, each URL counts as only 1 character towards the 280 character limit. After it’s submitted, the note is available for other users to read and rate. Birdwatch notes are public, and anyone can browse the Birdwatch website to see them.

**Ratings:** Users on Birdwatch can rate the helpfulness of notes from other users. These ratings are supposed to help to identify which notes are most helpful and to allow Birdwatch to raise the visibility of the context that is found most helpful by a wide range of contributors. Tweets with Birdwatch notes are highlighted with a Birdwatch icon that helps users to find them. Users can then click on the icon to read all notes others have written about that tweet and rate their quality (i. e., whether or not the Birdwatch note is helpful). Notes rated by the community to be particularly helpful then receive a currently rated helpful badge.

<sup>1</sup>The Birdwatch website is available via [birdwatch.twitter.com](https://birdwatch.twitter.com)

Source tweet	Birdwatch note		Helpful	
	Categorization	Text explanation	Yes	No
#1 – <b>Paul Elliott Johnson (@RhetoricPJ)</b> : “Rush Limbaugh had a regular radio segment where he would read off the names of gay people who died of AIDS and celebrate it and play horns and bells and stuff.”	<u>Misleading</u>	“These segments are fictitious and have never happened.”	3	21
#2 – <b>Alexandria Ocasio-Cortez (@AOC)</b> : “I am happy to work with Republicans on this issue where there’s common ground, but you almost had me murdered 3 weeks ago so you can sit this one out. Happy to work w/ almost any other GOP that aren’t trying to get me killed. In the meantime if you want to help, you can resign. [LINK]”	Misleading	“While I can understand AOC being worried she may have been hurt by the Capitol Protesters she cannot claim Ted Cruz tried to have her killed without providing proof. That’s dangerous misinformation by falsely accusing Ted Cruz of a serious crime.”	47	12
#3 – <b>Adam Kinzinger (@RepKinzinger)</b> : “The vast majority of Americans believe in universal background checks. As a gun owner myself, I firmly support the Second Amendment but I also believe we have to be willing to make some changes for the greater good. Read my full statement on #HR8 here: [LINK]”	Not misleading	“Universal background checks enjoy high levels of public support; a 2016 representative survey found 86% of registered voters in the United States supported the measure. [LINK]”	7	0

Table 1: Examples of Birdwatch notes and ratings. The column “Helpful” refers to the number of users who have responded “Yes” (“No”) to the rating question “Is this note helpful?” Wavy underlining indicates fact-checks that are clearly false (i.e., abuse of the fact-checking feature).

**Illustrative examples:** Tbl. 1 presents three examples of Birdwatch notes and helpfulness ratings. The table shows that many Birdwatch notes report tweets that refer to a political topic (e. g., voting rights). We also see that the text explanations differ regarding their informativeness. Some text explanations are longer and link to additional sources that are supposed to support the comments of the author. There are also users trying to abuse Birdwatch by falsely asserting that a tweet is misleading. In such situations, other users can down-vote the note through Birdwatch’s rating system. As an example, the (false) assertion in Birdwatch note #1 received 3 helpful and 21 unhelpful votes. In some cases, users also fact-check tweets for which veracity may be seen as being difficult to assess. For instance, Birdwatch note #2 addresses the source tweet from Democratic Congresswoman Alexandria Alexandria Ocasio-Cortez as a factual assertion that members of the GOP were trying to get her killed during the Capitol riots on January 6, 2021. However, one may also argue that the tweet can be seen a hyperbolic claim to emphasize that she believes that certain policy views or rhetorical stance encouraged the Capitol rioters. This is reflected in mixed helpfulness ratings: Birdwatch note #2 received 47 helpful and 12 unhelpful votes.

## Data

### Data Collection

We downloaded *all* Birdwatch notes and ratings between the introduction of the feature on January 23, 2021, and the end of July 2021 from the Birdwatch website. This comprehensive dataset contains a total number of 11,802 Birdwatch notes and 52,981 ratings. On Birdwatch, multiple users can

write Birdwatch notes for the same tweet. The average number of Birdwatch notes per fact-checked tweet is 1.31. Each Birdwatch note has a unique id (*noteId*) and each rating refers to a single Birdwatch note. We merged the ratings with the Birdwatch notes using this *noteId* field. The result is a single dataframe in which each Birdwatch note corresponds to one observation for which we know the number of helpful and unhelpful votes from the rating data.

### Key Variables

We now present the key variables that we extracted from the Birdwatch data. All of these variables will be empirically analyzed in the next sections:

- *Misleading*: A binary indicator of whether a tweet has been reported as being misleading by the author of the Birdwatch note (= 1; otherwise = 0).
- *Text explanation*: The user-entered text explanation (max 280 characters) to explain why a tweet is misleading or not misleading.
- *Trustworthy sources*: A binary indicator of whether the author of the Birdwatch note has responded “Yes” to the question “Did you link to sources you believe most people would consider trustworthy?” (= 1; otherwise = 0).
- *HVotes*: The number of other users who have responded “Yes” to the rating question “Is this note helpful?”
- *Votes*: The total number of ratings a Birdwatch note has received from other users (helpful & unhelpful).

### Content Characteristics

We use text mining methods to extract content characteristics from the text explanations of the Birdwatch notes:

- **Sentiment:** We calculate a sentiment score that measures the extent of positive vs. negative emotions in Birdwatch notes. Our computation follows a dictionary-based approach as in Vosoughi, Roy, and Aral (2018). Here we use the NRC emotion lexicon (Mohammad and Turney 2013), which classifies English words into positive and negative emotions. The fraction of words in the text explanations related to positive and negative emotions is then aggregated and averaged to create a vector of positive and negative emotion weights that sum to one. The sentiment score is then defined as the difference between positive and negative emotion scores.
- **Text complexity:** We calculate a text complexity score, specifically, the Gunning-Fog index (Gunning 1968). This index estimates the years of formal education necessary for a person to understand a text upon reading it for the first time:  $0.4 \times (ASL + 100 \times n_{wsy \geq 3} / n_w)$ , where  $ASL$  is the average sentence length (number of words),  $n_w$  is the total number of words, and  $n_{wsy \geq 3}$  is the number of words with three syllables or more. A higher value thus indicates greater complexity.
- **Word count:** We determine the length of the text explanations in Birdwatch notes as given by the number of words.

Our text mining pipeline is implemented in R 4.0.2 using the packages `quanteda` (Benoit et al. 2018) in Version 2.0.1 and `sentimentr` (Rinker 2019) in Version 2.7.1.

## Twitter Historical API

Each Birdwatch note addresses a single tweet, i. e., the tweet that has been categorized as being misleading or not misleading by the author of the Birdwatch note. We used the Twitter historical API to map the `tweetID` referenced in each Birdwatch note to the source tweet and collected the following information for the author of each source tweet:

- **Account name:** The name of the Twitter account for which the Birdwatch note has been reported.
- **Followers:** The number of followers, i. e., the number of accounts that follow the author of the source tweet.
- **Followees:** The number of followees, i. e., the number of accounts whom the author of the source tweet follows.
- **Account age:** The age of the author of the source tweet’s account (in years).
- **Verified:** A binary dummy indicating whether the account of the source tweet has been officially verified by Twitter (= 1; otherwise = 0).

## Empirical Analysis

### Analysis of Birdwatch Notes (RQ1 & RQ2)

In this section, we analyze how users categorize misleading and not misleading tweets in Birdwatch notes. We also explore the reasons because of which Birdwatch users report tweets and how Birdwatch notes for misleading vs. not misleading differ in terms of their content characteristics.

**Categorization of tweets in Birdwatch notes:** Fig. 2 shows that users file a larger number of Birdwatch notes for

misleading than for not misleading tweets. Out of all Birdwatch notes, 90 % refer to misleading tweets, whereas the remaining 10 % refer to not misleading tweets. Fig. 2 further suggests that authors of Birdwatch notes are more likely to report that they have linked to trustworthy sources for misleading (75 %) than for not misleading tweets (64 %).




Figure 2: Number of users who responded “Yes” to the question “Did you link to sources you believe most people would consider trustworthy?”

**Why do users report misleading tweets?** Authors of Birdwatch notes also need to answer a checkbox question (multiple selections possible) on why they perceive a tweet as being misleading. Fig. 3 shows that misleading tweets are primarily reported because of factual errors (32 %), lack of important context (30 %), or because they treat unverified claims as facts (26 %). Birdwatch notes reporting outdated information (5 %), satire (3 %), or manipulated media (2 %) are relatively rare. Only 3 % of Birdwatch notes are reported because of other reasons.




Figure 3: Number of Birdwatch notes per checkbox answer option in response to the question “Why do you believe this tweet may be misleading?”

**Why do users report not misleading tweets?** Fig. 4 shows that not misleading tweets are primarily reported because they are perceived as factually correct (58 %), clearly refer to personal opinion (19 %), or satire (12 %). Birdwatch users only rarely report tweets containing information that was correct at the time of writing but is now outdated (1 %). 10 % of tweets are reported because of other reasons.

**Content characteristics:** The complementary cumulative distribution functions (CCDFs) in Fig. 5 visualize how Birdwatch notes for misleading vs. not misleading tweets differ in terms of their content characteristics. We find that Birdwatch notes reporting misleading tweets use similarly complex language but embed a higher proportion of negative emotions than tweets reported as being not misleading. For tweets reported as being misleading, authors of Birdwatch notes also tend to write longer text explanations.






Figure 4: Number of Birdwatch notes per checkbox answer option in response to the question “Why do you believe this tweet is not misleading?”

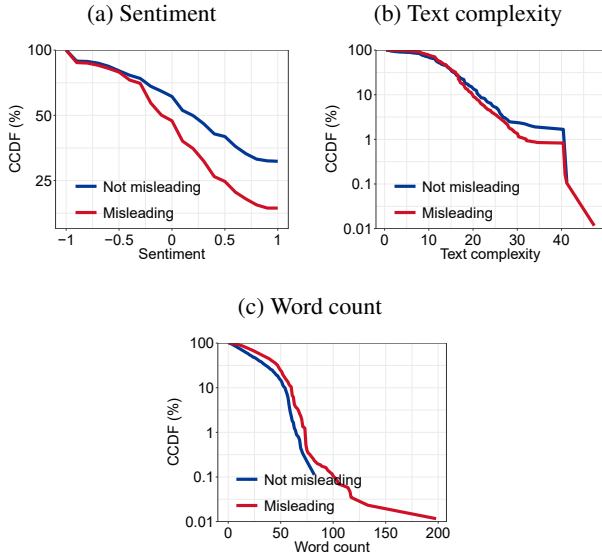


Figure 5: CCDFs for (a) sentiment, (b) text complexity, and (c) word count in text explanations of Birdwatch notes.

### Analysis of Source Tweets (RQ3)

We now explore whether tweets from Twitter accounts with certain characteristics (e. g., politicians, accounts with many followers) are more likely to be fact-checked on Birdwatch.

**Most reported accounts:** Tbl. 2 reports the names of the ten Twitter accounts that have gathered the most Birdwatch notes. The table also reports the number of unique users that have filed Birdwatch notes per account. We find that many of the most-reported accounts belong to influential politicians from both sides of the political spectrum in the U. S. – for which Birdwatch users overwhelmingly report misinformation. For instance, 160 Birdwatch notes have been reported for the account of Republican U. S. House Representative Marjorie Taylor Greene (@mtgreenee), out of which 97% have been categorized as being misleading tweets. Similarly, 160 Birdwatch notes have been reported for Democratic Congresswoman Alexandria Ocasio-Cortez (AOC), out of which 90% have been categorized as being misleading tweets. The highest number of Birdwatch notes has been filed for @Quakeprediction, an account that regu-

larly publishes earthquake predictions. Notably, this account has only been fact-checked by three different users that have filed an extensive number of fact-checks. We further observe many fact-checks for newspapers (@thehill) and other prominent public figures such as NYT best-selling author Candace Owens (@RealCandaceO).

Account name	#Notes (% misleading)	#Users
EarthquakePrediction (@Quakeprediction)	257 (100%)	3
Marjorie Taylor Greene (@mtgreenee)	180 (97%)	93
Alexandria Ocasio-Cortez (@AOC)	160 (90%)	105
Lauren Boebert (@laurenboebert)	108 (92%)	62
Alex Berenson (@AlexBerenson)	78 (97%)	44
President Biden (@POTUS) <sup>†</sup>	74 (81%)	58
The Hill (@thehill)	66 (95%)	50
Candace Owens (@RealCandaceO)	49 (90%)	46
Japan Earthquakes (@earthquakejapan)	44 (100%)	1
Aaron Rupar (@atrupar)	39 (85%)	27

<sup>†</sup>The presidential Twitter account @POTUS was held by Joe Biden throughout the entire study period. The account has been reset after the transition to the new administration on January 20, 2021.

Table 2: Top-10 Twitter accounts with the highest number of Birdwatch notes.

Next, we explore Twitter accounts with the highest share of tweets reported as being misleading (Tbl. 3). For this purpose, we employed the Twitter historical API to obtain the number of tweets posted by each fact-checked account (i. e., the user timelines) since the introduction of Birdwatch. To identify accounts that received interest by many fact-checkers, we focus on accounts that have been fact-checked by at least 20 unique users. We again find that Birdwatch users have pronounced interest in debunking posts from politicians. The highest share of misleading tweets is found for the account of Alexandria Ocasio-Cortez, for which approximately 6% of all posted tweets have been reported as being misleading. Note that Marjorie Taylor Greene appears in the top-10 list with both her private (@mtgreenee) and her official governmental account (@RepMTG).

Account name	#Posts (% Misleading)	#Users
Alexandria Ocasio-Cortez (@AOC)	640 (5.6%)	105
Rep. Marjorie Taylor Greene (@RepMTG)	303 (5.6%)	22
Marjorie Taylor Greene (@mtgreenee)	2,182 (4.9%)	93
Lauren Boebert (@laurenboebert)	1,486 (4.3%)	62
Candace Owens (@RealCandaceO)	576 (3.8%)	46
President Biden (@POTUS)	1,176 (3.3%)	58
Alex Berenson (@AlexBerenson) <sup>†</sup>	2,699 (2.0%)	44
Rep. Jim Jordan (@Jim Jordan)	1,273 (1.8%)	25
Cori Bush (@CoriBush)	414 (1.7%)	27
Charlie Kirk (@charliekirk11)	1,051 (1.4%)	21

<sup>†</sup>Account has been suspended by Twitter.

Table 3: Top-10 Twitter accounts with the highest share of misleading tweets. Here we focus on accounts that have been fact-checked by at least 20 different users.

**Account characteristics:** Fig. 6 plots the CCDFs for different characteristics of Twitter accounts, namely (a) the number of followers, (b) the number of followees, and (c) the account age (in years). The plots show that tweets reported as being misleading tend to have a lower number of followers and a lower number of followees. We observe no significant differences with regard to the account age. As a further analysis, we investigated differences across verified vs. not verified accounts. Here we find that 54 % of all misleading tweets were posted by verified accounts, whereas this number is 61 % for not misleading tweets.

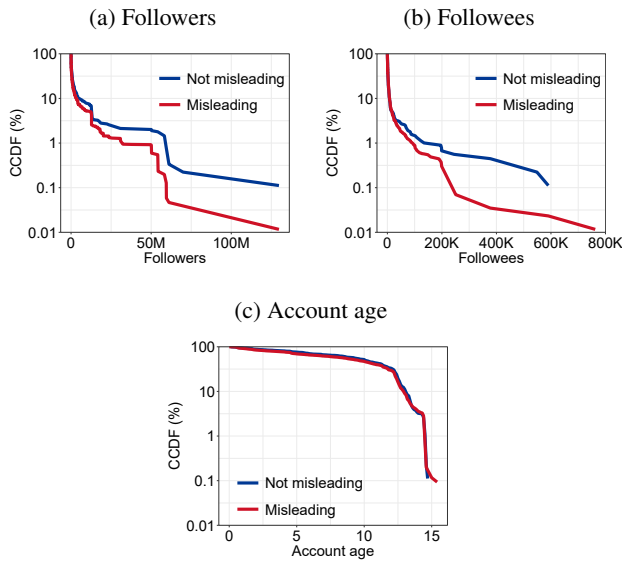


Figure 6: CCDFs for (a) followers, (b) followees, and (c) account age (in years).

### Analysis of Ratings (RQ4)

**Helpfulness:** Fig. 7 plots the CCDFs for the helpfulness ratio ( $HVotes/Votes$ ) and the total number of helpful and unhelpful votes ( $Votes$ ). We find that Birdwatch notes reporting misleading tweets tend to have a higher helpfulness ratio but receive a lower number of total votes.




Figure 7: CCDFs for (a) helpfulness ratio and (b) total votes.

**Why do users find Birdwatch notes helpful?** Users rating a Birdwatch note helpful need to answer a checkbox

question (multiple selections possible) on why they perceive it as being helpful. Fig. 8 shows that users find that Birdwatch notes are particularly helpful if they are (i) informative (24 %), (ii) clear (24 %), and (iii) provide good sources (21 %). To a lesser extent, users also value Birdwatch notes that provide unique/informative context (14 %) and are empathetic (11 %). Only 6 % of Birdwatch notes are perceived as helpful because of other reasons.

**Why do users find Birdwatch notes unhelpful?** Fig. 9 shows that users find Birdwatch notes unhelpful (i) if sources are missing or unreliable (19 %), (ii) if there is opinion speculation or bias (19 %), (iii) if key points are missing (18 %), (iv) if it is argumentative or inflammatory (13 %), or (v) if it is incorrect (10 %). In some cases, users also perceive a Birdwatch note as unhelpful because it is off-topic (5 %) or hard to understand (4 %). Only few Birdwatch notes are perceived as unhelpful because of spam harassment (3 %), outdated information (2 %), irrelevant sources (2 %), or other reasons (5 %).




Figure 8: Number of ratings per checkbox answer option in response to the prompt “What about this note was helpful to you?”




Figure 9: Number of ratings per checkbox answer option in response to the question “Help us understand why this note was unhelpful.”

### Regression Model for Helpfulness (RQ4 & RQ5)

We now address the question of “what makes a Birdwatch note helpful?” To address this question, we employ regression analyses using helpfulness as the dependent variable and (i) the user categorization, (ii) the content characteristics of the text explanation, (iii) the social influence of the source tweet as explanatory variables.

**Model specification:** Following previous research modeling helpfulness (e.g., Lutz, Pröllochs, and Neumann 2022),



we model the number of helpful votes,  $HVotes$ , as a binomial variable with probability parameter  $\theta$  and  $Votes$  trials:

$$\begin{aligned} \text{logit}(\theta) = & \beta_0 + \underbrace{\beta_1 \text{Misleading} + \beta_2 \text{TrustworthySources}}_{\text{User categorization}} \quad (1) \\ & + \underbrace{\beta_3 \text{TextComplexity} + \beta_4 \text{Sentiment} + \beta_5 \text{WordCount}}_{\text{Text explanation}} \\ & + \underbrace{\beta_6 \text{AccountAge} + \beta_7 \text{Followers} + \beta_8 \text{Followees} + \beta_9 \text{Verified}}_{\text{Source tweet}} + \varepsilon, \end{aligned}$$

$$HVotes \sim \text{Binomial}[Votes, \theta], \quad (2)$$

with intercept  $\beta_0$  and error term  $\varepsilon$ . We estimate Eq. 1 and Eq. 2 using maximum likelihood estimation and generalized linear models. To facilitate the interpretability of our findings, we  $z$ -standardize all variables, so that we can compare the effects of regression coefficients on the dependent variable measured in standard deviations.

We use the same set of explanatory variables as in Eq. 1 to model the total number of votes (helpful and unhelpful). The dependent variable is thus given by  $Votes$ . We estimate this model via a negative binomial regression with log-transformation. The reason is that the number of votes denotes count data with overdispersion (i. e., variance larger than the mean).

**Coefficient estimates:** The parameter estimates in Fig. 10 show that Birdwatch notes reporting misleading tweets are significantly more helpful. The coefficient for *Misleading* is 0.465 ( $p < 0.001$ ), which implies that the odds of Birdwatch notes reporting misleading tweets are  $e^{0.465} \approx 1.59$  times the odds for Birdwatch notes reporting not misleading tweets. Similarly, we find that the odds for Birdwatch notes providing trustworthy sources are 2.01 times the odds for Birdwatch notes that do not provide trustworthy sources (coef: 0.698;  $p < 0.001$ ). We also find that content characteristics of Birdwatch notes are important determinants regarding their helpfulness: the coefficients for *Text complexity*, *Word count*, and *Sentiment* are positive and statistically significant ( $p < 0.001$ ). Hence, Birdwatch notes are estimated to be more helpful if they embed positive sentiment, and if they use more words and more complex language. For a one standard deviation increase in the explanatory variable, the estimated odds of a helpful vote increase by a factor of 1.15 for *Text complexity*, 1.10 for *Word count*, and 1.06 for *Sentiment*. We also find that the social influence of the account that has posted the source tweet plays an important role regarding the helpfulness of Birdwatch notes. Here the largest effect sizes are estimated for *Verified* and *Followers* ( $p < 0.001$ ). The odds for Birdwatch notes reporting tweets from verified accounts are 0.76 times the odds for unverified accounts. A one standard deviation increase in the number of followers decreases the odds of a helpful vote by a factor of 0.81. In sum, there is a lower level of consensus for verified and high-follower accounts.

We observe a different pattern for the negative binomial model with the total number of votes (helpful and unhelpful) as the dependent variable. Here we find a negative and statistically significant coefficient for *Account age* ( $p < 0.001$ ) and positive and statistically significant coefficients for *Verified* and *Followers* ( $p < 0.001$ ). A one standard deviation

increase in the *Account age* is expected to decrease the total number of votes by a factor of 0.84. A one standard deviation increase in the number of followers is expected to increase the total number of votes by a factor of 1.81. Birdwatch notes reporting tweets from verified accounts are expected to receive 2.30 times more votes. Altogether, we find that the total number of votes is strongly associated with the social influence of the source tweet rather than the characteristics of the Birdwatch note itself. In contrast, the odds of a Birdwatch note to be helpful are strongly associated with both the social influence of the source tweet and the characteristics of the Birdwatch note. Birdwatch notes reported for accounts with high social influence receive more total votes but are less likely to be perceived as being helpful.

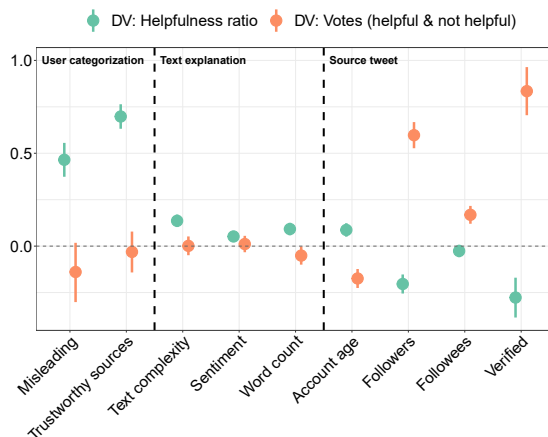


Figure 10: Regression results for helpfulness ratio and total votes as dependent variables (DV). Reported are standardized parameter estimates and 99% confidence intervals.

**Analysis for subcategories of helpfulness:** We repeat our regression analysis for different subcategories of helpfulness. For this purpose, we estimate the effect of the explanatory variables on the frequency of the responses of Birdwatch users to the question “What about this note was helpful/unhelpful to you?” For brevity, we omit the coefficient estimates and instead briefly describe the main findings. In short, we find that Birdwatch notes reporting misleading tweets are more informative but not necessarily more empathetic. We further observe that Birdwatch notes for misleading tweets are less likely spam harassment, opinion speculation or hard to understand. We again find that linking to trustworthy sources is crucial regarding helpfulness. Unsurprisingly, the largest effect size is observed for *Trustworthy sources* on the dependent variable *Good sources*. However, linking trustworthy sources also makes it more likely that a Birdwatch note is perceived to provide unique context, to be informative, to be empathetic, and to be clear. We find that Birdwatch notes that are longer, more complex, and embed a more positive sentiment are perceived as more helpful across most subcategories of helpfulness. Also, more positive sentiment is less likely to be perceived as argumentative. Our analysis further suggests reasons be-

cause of which fact checks from verified and high-follower accounts are perceived as less helpful: Birdwatch notes for these accounts are more likely to be seen as being incorrect, argumentative, opinion speculation, and spam harassment.

**Model checks:** (1) We checked that variance inflation factors as an indicator of multicollinearity were below five. (2) We controlled for non-linear relationships, user-specific effects, and month-level time effects. (3) We estimated separate regressions for misleading vs. not misleading tweets. In all cases, our results are robust and consistently support our findings.

## Discussion

In this work, we provide a holistic analysis of the new Birdwatch pilot on Twitter. Our empirical findings contribute to research into misinformation and crowdsourced fact-checking. Complementary to experimental studies focusing on the question of whether crowds are *able* to accurately assess social media content (e.g., Pennycook and Rand 2019a), this work sheds light on how users *interact* with community-based fact-checking systems.

**Implications:** Our work has direct implications for the Birdwatch platform and future attempts to implement community-based approaches to combat misinformation on social media. We find that Twitter has developed a fact-checking platform on which users perceive a relatively high share of community-created fact checks (Birdwatch notes) as being informative, clear and helpful. These findings are encouraging considering that earlier crowd-based fact-checking initiatives (e.g., *TruthSquad*, *WikiTribune*) were largely unsuccessful to produce high-quality fact-checks at scale (Bhuiyan et al. 2020). A potential reason may be that Birdwatch differs in several important ways: (i) it is directly integrated into the Twitter platform and thus has relatively low entry barriers for fact-checkers; (ii) it features a relatively sophisticated rating system to promote helpful fact-checks; (iii) it is a purely community-based approach which may reduce the problem that many users distrust professional fact-checkers (Poynter 2019); (iv) it encourages users to provide context and link to trustworthy sources, which is known to enhance credibility (Zhang et al. 2018).

Despite showing promising potential, our findings suggest that community-based fact-checking faces challenges with regards to social media content from influential user accounts. On Birdwatch, a large share of fact-checks is carried out for tweets from influential politicians from both sides of the political spectrum in the U.S.—for which users overwhelmingly report misinformation. This indicates that users have strong interest in debunking political misinformation. We can only speculate whether tweets from these accounts are more often misleading; or rather that the fact-checking community is more likely to report posts that do not align with their political ideology. What we do find, however, is that the level of consensus between users is significantly lower if misinformation is reported for tweets from users with influential accounts. Our analysis of subcategories of helpfulness further reveals that these fact-checks are particularly likely to be perceived as being argumentative, incorrect or opinion speculation. In line with earlier research

(Soares, Recuero, and Zago 2018; Barberá et al. 2015), this suggests that influential user accounts (e.g., politicians) play a key role in the formation of fragmented social networks and that there is a high level of polarization for the content they produce. Polarized social media users often have vastly different interpretations of facts or even entirely different sets of acknowledged facts (Ojala et al. 2021). For crowd-based fact-checking platforms, this implies that they need to ensure independence and high diversity among fact-checkers, such that negative effects are eventually mitigated (Epstein, Pennycook, and Rand 2020).

**Fact-checking guidelines:** Our findings may help to formulate guidelines for users on how to write more helpful fact checks. We find that it is crucial that users provide *context* in their text explanations. Users should thus be encouraged to link trustworthy sources in their fact-checks. Furthermore, we show that content characteristics matter. Community-created fact-checks are estimated to be more helpful if they embed more positive vs. negative emotions. Users should thus avoid employing inflammatory language. We also find that community-created fact-checks that use more words and complex language are perceived as more helpful. It should thus be encouraged that users provide profound explanations and exhaust the full character limit to explain why they perceive a tweet as being misleading or not misleading.

**Limitations and directions for future research:** This work has a number of limitations, which can fuel future research as follows. First, future research should conduct analyses to better understand which groups of users engage in fact-checking and the accuracy of the community-created fact checks. While previous works suggest that collective wisdom can be more accurate than individual’s assessments and even experts (Bhuiyan et al. 2020), there are situations in which the crowd may perform worse. For example, it is an ongoing discussion whether crowds become more or less accurate in sequential decision making problems (Frey and van de Rijt 2021). More research is thus necessary to better understand the conditions under which the wisdom of crowds can be unlocked for fact-checking. Second, it will be of utmost importance to understand the effect the Birdwatch feature has on the user behavior on the overall social media platform (i.e., Twitter). It has to be ensured that Birdwatch does not yield adverse side effects. Previous experience of Facebook with “flags” suggests that adding warning labels to some posts can make all other (unverified) posts seem more credible (BBC 2017). Future research should seek to understand the effects of community-created fact checks on the diffusion of both the fact-checked tweet itself as well as other misleading information circulating on Twitter.

Finally, our inferences are limited to community-based fact-checking on Twitter’s Birdwatch platform. Further research is thus necessary to analyze whether the observed patterns are generalizable to other crowd-based fact-checking platforms. Also, the characteristics of the restricted set of users participating in the current Birdwatch pilot phase may differ from the overall user base on social media platforms. For example, one may expect increased diversity for wider audiences, for which the wisdom of crowds literature suggests that crowds can perform better (Becker, Brackbill, and

Centola 2017). Notwithstanding these limitations, we believe that Birdwatch provides a unique opportunity to empirically analyze how users interact with a novel community-based fact-checking system, and that observing and understanding how users interact with such a system is the first step towards its rollout to wider audiences.

## Conclusion

Twitter recently launched Birdwatch, a community-driven approach to identify misinformation on social media by harnessing the wisdom of crowds. This study is the first to empirically analyze how users interact with this new feature. We find that encouraging users to provide context (e. g., by linking trustworthy sources) and to avoid the use of inflammatory language is crucial for users to perceive community-based fact checks as helpful. However, despite showing promising potential, our analysis also suggests that Birdwatch's community-driven approach faces challenges concerning opinion speculation and polarization among the user base – in particular with regards to influential user accounts. Our findings are relevant both for the Birdwatch platform and future attempts to implement community-based approaches to combat misinformation on social media.

## Acknowledgments

This study was supported by a research grant from the German Research Foundation (DFG grant 492310022).

## References

- Allen, J.; Arechar, A. A.; Pennycook, G.; and Rand, D. G. 2021. Scaling up fact-checking using the wisdom of crowds. *Science Advances*, 7(36): eabf4393.
- Allen, J.; Howland, B.; Mobius, M.; Rothschild, D.; and Watts, D. J. 2020. Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14): eaay3539.
- Aral, S.; and Eckles, D. 2019. Protecting elections from social media manipulation. *Science*, 365(6456): 858–861.
- Bakabar, M. 2018. Crowdsourced factchecking. <https://fullfact.org/blog/2018/may/crowdsourced-factchecking/>. Accessed: 2022-04-02.
- Bakshy, E.; Messing, S.; and Adamic, L. A. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239): 1130–1132.
- Barberá, P.; Jost, J. T.; Nagler, J.; Tucker, J. A.; and Bonneau, R. 2015. Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, 26(10): 1531–1542.
- BBC. 2017. Facebook ditches fake news warning flag. <https://www.bbc.com/news/technology-42438750>. Accessed: 2022-04-02.
- Becker, J.; Brackbill, D.; and Centola, D. 2017. Network dynamics of social influence in the wisdom of crowds. *PNAS*, 114(26): 5070–5076.
- Benoit, K.; Watanabe, K.; Wang, H.; Nulty, P.; Obeng, A.; Müller, S.; and Matsuo, A. 2018. quanteda: An R package for the quantitative analysis of textual data. *Journal of Open Source Software*, 3(30): 774.
- Bhuiyan, M. M.; Zhang, A. X.; Sehat, C. M.; and Mitra, T. 2020. Investigating differences in crowdsourced news credibility assessment: Raters, tasks, and expert criteria. *Proceedings of the ACM on Human-Computer Interaction*, 4: 1–26.
- Conover, M. D.; Ratkiewicz, J.; Francisco, M.; Gonçalves, B.; Menczer, F.; and Flammini, A. 2011. Political polarization on Twitter. In *ICWSM*.
- Del Vicario, M.; Bessi, A.; Zollo, F.; Petroni, F.; Scala, A.; Caldarelli, G.; Stanley, H. E.; and Quattrociocchi, W. 2016. The spreading of misinformation online. *PNAS*, 113(3): 554–559.
- Ducci, F.; Kraus, M.; and Feuerriegel, S. 2020. Cascade-LSTM: A tree-structured neural classifier for detecting misinformation cascades. In *KDD*.
- Ecker, U. K.; Lewandowsky, S.; and Tang, D. T. 2010. Explicit warnings reduce but do not eliminate the continued influence of misinformation. *Memory & Cognition*, 38(8): 1087–1100.
- Epstein, Z.; Pennycook, G.; and Rand, D. 2020. Will the crowd game the algorithm? Using layperson judgments to combat misinformation on social media by downranking distrusted sources. In *CHI*.
- Florin, F. 2010. Crowdsourced fact-checking? What we learned from Truthsquad. <http://mediashift.org/2010/11/crowdsourced-fact-checking-what-we-learned-from-truthsquad320/>. Accessed: 2022-04-02.
- Frey, V.; and van de Rijt, A. 2021. Social influence undermines the wisdom of the crowd in sequential decision making. *Management Science*, 67(7): 4273–4286.
- Friggeri, A.; Adamic, L. A.; Eckles, D.; and Cheng, J. 2014. Rumor cascades. In *ICWSM*.
- Geeng, C.; Yee, S.; and Roesner, F. 2020. Fake News on Facebook and Twitter: Investigating How People (Don't) Investigate. In *CHI*.
- Godel, W.; Sanderson, Z.; Aslett, K.; Nagler, J.; Bonneau, R.; Persily, N.; and Tucker, J. A. 2021. Moderating with the mob: evaluating the efficacy of real-time crowdsourced fact-checking. *Journal of Online Trust and Safety*, 1(1): 1–36.
- Gunning, R. 1968. *The technique of clear writing*. New York: McGraw-Hill.
- Hassan, N.; Arslan, F.; Li, C.; and Tremayne, M. 2017. Toward automated fact-checking: Detecting check-worthy factual claims by claimbuster. In *KDD*.
- Kahan, D. M. 2017. Misconceptions, misinformation, and the logic of identity-protective cognition. SSRN:2973067.
- Kim, A.; and Dennis, A. R. 2019. Says who? The effects of presentation format and source rating on fake news in social media. *MIS Quarterly*, 43(3): 1025–1039.
- Kwon, S.; and Cha, M. 2014. Modeling bursty temporal pattern of rumors. In *ICWSM*.
- Lazer, D. M. J.; Baum, M. A.; Benkler, Y.; Berinsky, A. J.; Greenhill, K. M.; Menczer, F.; Metzger, M. J.; Nyhan,

- B.; Pennycook, G.; Rothschild, D.; Schudson, M.; Sloman, S. A.; Sunstein, C. R.; Thorson, E. A.; Watts, D. J.; and Zittrain, J. L. 2018. The science of fake news. *Science*, 359(6380): 1094–1096.
- Levy, R. 2021. Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3): 831–70.
- Luca, M.; and Zervas, G. 2016. Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science*, 62(12): 3412–3427.
- Lutz, B.; Pröllochs, N.; and Neumann, D. 2022. Are longer reviews always more helpful? Disentangling the interplay between review length and line of argumentation. *Journal of Business Research*, 144: 888–901.
- Ma, J.; Gao, W.; Mitra, P.; Kwon, S.; Jansen, B. J.; Wong, K.-F.; and Cha, M. 2016. Detecting rumors from microblogs with recurrent neural networks. In *ICJAI*.
- McPherson, M.; Smith-Lovin, L.; and Cook, J. M. 2001. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27: 415–444.
- Micallef, N.; He, B.; Kumar, S.; Ahamad, M.; and Memon, N. 2020. The role of the crowd in countering misinformation: A case study of the COVID-19 infodemic. In *International Conference on Big Data*.
- Mohammad, S. M.; and Turney, P. D. 2013. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3): 436–465.
- Moravec, P. L.; Minas, R. K.; and Dennis, A. 2019. Fake news on social media: People believe what they want to believe when it makes no sense at all. *MIS Quarterly*, 43(4): 1343–1360.
- Okoli, C.; Mehdi, M.; Mesgari, M.; Nielsen, F. Å.; and Lanamäki, A. 2014. Wikipedia in the eyes of its beholders: A systematic review of scholarly research on Wikipedia readers and readership. *Journal of the Association for Information Science and Technology*, 65(12): 2381–2403.
- O’Riordan, S.; Kiely, G.; Emerson, B.; and Feller, J. 2019. Do you have a source for that? Understanding the challenges of collaborative evidence-based journalism. In *OpenSym*.
- Otala, J.; Kurtic, G.; Grasso, I.; Liu, Y.; Matthews, J.; and Madraki, G. 2021. Political polarization and platform migration: A Study of Parler and Twitter usage by United States of America Congress members. In *WWW Companion*.
- Pennycook, G.; Cannon, T. D.; and Rand, D. G. 2018. Prior exposure increases perceived accuracy of fake news. *Journal of Experimental Psychology: General*, 147(12): 1865–1880.
- Pennycook, G.; Epstein, Z.; Moseleh, M.; Arechar, A. A.; Eckles, D.; and Rand, D. G. 2021. Shifting attention to accuracy can reduce misinformation online. *Nature*, 592(7855): 590–595.
- Pennycook, G.; and Rand, D. G. 2019a. Fighting misinformation on social media using crowdsourced judgments of news source quality. *PNAS*, 116(7): 2521–2526.
- Pennycook, G.; and Rand, D. G. 2019b. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188: 39–50.
- Pew Research Center. 2016. News use across social media platforms 2016. <https://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/>. Accessed: 2022-04-02.
- Poynter. 2019. Most Republicans don’t trust fact-checkers, and most Americans don’t trust the media. <https://www.poynter.org/ifcn/2019/most-republicans-dont-trust-fact-checkers-and-most-americans-dont-trust-the-media/>. Accessed: 2022-04-02.
- Pröllochs, N.; Bär, D.; and Feuerriegel, S. 2021a. Emotions explain differences in the diffusion of true vs. false social media rumors. *Scientific Reports*, 11: 22721.
- Pröllochs, N.; Bär, D.; and Feuerriegel, S. 2021b. Emotions in online rumor diffusion. *EPJ Data Science*, 10(1): 51.
- Qazvinian, V.; Rosengren, E.; Radev, D. R.; and Mei, Q. 2011. Rumor has it: Identifying misinformation in microblogs. In *EMNLP*.
- Rinker, T. W. 2019. *sentimentr: Calculate Text Polarity Sentiment*. Buffalo, New York. Version 2.7.1.
- Shao, C.; Ciampaglia, G. L.; Flammini, A.; and Menczer, F. 2016. Hoaxy: A platform for tracking online misinformation. In *WWW Companion*.
- Soares, F. B.; Recuero, R.; and Zago, G. 2018. Influencers in polarized political networks on Twitter. In *SMSociety*.
- Solovev, K.; and Pröllochs, N. 2022a. Hate speech in the political discourse on social media: Disparities across parties, gender, and ethnicity. In *WWW*.
- Solovev, K.; and Pröllochs, N. 2022b. Moral emotions shape the virality of COVID-19 misinformation on social media. In *WWW*.
- Starbird, K. 2017. Examining the alternative media ecosystem through the production of alternative narratives of mass shooting events on Twitter. In *ICWSM*.
- Vo, N.; and Lee, K. 2018. The rise of guardians: Fact-checking url recommendation to combat fake news. In *SIGIR*.
- Vosoughi, S.; Mohsenvand, M. N.; and Roy, D. 2017. Rumor gauge: Predicting the veracity of rumors on Twitter. *ACM Transactions on Knowledge Discovery from Data*, 11(4): 1–36.
- Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true and false news online. *Science*, 359(6380): 1146–1151.
- Woolley, A. W.; Chabris, C. F.; Pentland, A.; Hashmi, N.; and Malone, T. W. 2010. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004): 686–688.
- Wu, L.; Morstatter, F.; Carley, K. M.; and Liu, H. 2019. Misinformation in social media: definition, manipulation, and detection. *SIGKDD Explorations Newsletter*, 21(2): 80–90.
- Zhang, A. X.; Ranganathan, A.; Metz, S. E.; Appling, S.; Sehat, C. M.; Gilmore, N.; Adams, N. B.; Vincent, E.; Lee, J.; Robbins, M.; et al. 2018. A structured response to misinformation: Defining and annotating credibility indicators in news articles. In *WWW Companion*.