

# Justice in Misinformation Detection Systems: An Analysis of Algorithms, Stakeholders, and Potential Harms

Terrence Neumann

University of Texas at Austin  
Austin, TX, USA

Terrence.Neumann@mcombs.utexas.edu

Maria De-Arteaga

University of Texas at Austin  
Austin, TX, USA

dearteaga@mcombs.utexas.edu

Sina Fazelpour

Northeastern University  
Boston, MA, USA

s.fazel-pour@northeastern.edu

## ABSTRACT

Faced with the scale and surge of misinformation on social media, many platforms and fact-checking organizations have turned to algorithms for automating key parts of misinformation detection pipelines. While offering a promising solution to the challenge of scale, the ethical and societal risks associated with algorithmic misinformation detection are not well-understood. In this paper, we employ and extend upon the notion of *informational justice* to develop a framework for explicating issues of justice relating to representation, participation, distribution of benefits and burdens, and credibility in the misinformation detection pipeline. Drawing on the framework: (1) we show how injustices materialize for stakeholders across three algorithmic stages in the pipeline; (2) we suggest empirical measures for assessing these injustices; and (3) we identify potential sources of these harms. This framework should help researchers, policymakers, and practitioners reason about potential harms or risks associated with these algorithms and provide conceptual guidance for the design of algorithmic fairness audits in this domain.

## KEYWORDS

algorithmic fairness, justice, misinformation detection, machine learning, informational justice

### ACM Reference Format:

Terrence Neumann, Maria De-Arteaga, and Sina Fazelpour. 2022. Justice in Misinformation Detection Systems: An Analysis of Algorithms, Stakeholders, and Potential Harms. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, June 21–24, 2022, Seoul, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3531146.3533205>

## 1 INTRODUCTION

Misinformation has unfortunately become ubiquitous in our day-to-day life. The surge in the spread of misinformation on social networks has raised deep concerns regarding the well-being of individuals, groups, and democratic societies. Faced with the scale of the challenge, many platforms have turned to algorithms to automate content moderation decisions, and fact-checking organizations are increasingly adopting algorithmic tools for support.

Algorithmic content moderation appears to be widely adopted by some platforms. For instance, Facebook reportedly labeled 180 million messages as potentially misleading during the 2020 election season, while Twitter says it labeled 300,000 during the same period [50]. Further, Facebook claims to have taken 1.8 billion actions against fake accounts violating internal policies from July to September, 2021 [27].

However, there is a lack of attention to the ethical risks that algorithmic biases, similar to those found in systems used to assist with hiring, policing, and medical diagnostic decisions, may present in automated misinformation detection systems as well. Specifically, there has been no systematic study of the senses and sources of injustice that can plague automated misinformation detection systems. This gap is problematic not only because of what is at stake in how we devise and evaluate sociotechnical responses to misinformation; it is troublesome particularly because the much-discussed frameworks of justice that are suitable for allocation settings are not sufficient to capture all dimensions of justice in information ecosystems.

In this paper, we develop a framework for characterizing, detecting, and diagnosing the sources of injustices that can affect the performance of large-scale automated misinformation detection systems. We begin with discussing relevant existing work in Section 2. In Section 3, we introduce our framework by building on the notion of *informational justice*. Developed by Mathiesen [54], this multifaceted notion identifies three key dimensions of justice—participatory, recognitional, and distributive—that are pertinent to three types of stakeholders implicated in informational items—sources, subjects, and seekers of information, respectively. We extend this framework, to capture the additional stakeholders involved in misinformation detection pipelines and the further senses of justice that corresponds to them. In Section 4, we put this framework to use by situating it in the misinformation detection pipeline. Focusing on each stage of the pipeline, we explain why injustices of a specific sense might emerge in that stage, how we might empirically track their presence and trace their potential sources.

Throughout, we show how the framework provides practitioners, policy-makers, and platform executives with a robust tool for investigating the complex issues of justice, involving a myriad of diverse stakeholders, associated with automatic content moderation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9352-2/22/06...\$15.00

<https://doi.org/10.1145/3531146.3533205>

## 2 BACKGROUND AND RELATED WORK

### 2.1 Machine Learning and Algorithmic Injustices

The study of injustices that may result from the deployment of machine learning systems has largely focused on the risks of algorithmic bias [10], paying special attention to classification tasks that may inform the allocations of benefits and burdens [28, 58]. In such contexts, researchers have emphasized the risks of compounding inequalities [25] and historical injustices [40] in domains such as criminal justice [3], human resources [18], and healthcare [62]. While these relate primarily to allocative harms [9], representational harms have also received attention in the information retrieval context, with special attention devoted to stereotyping [13, 57].

Most closely related to our work is research on algorithmic fairness in the context of hate speech detection, as it constitutes another type of harmful online content [8]. Research in this space has empirically shown that automated hate speech detection systems may incorrectly label African American English tweets as toxic at disproportionate rates [65], and similar patterns of errors are observed for statements discussing identity of belonging to minoritized groups, such as posts saying “I am gay” [19]. This line of research has also provided critical perspectives of the datasets used to train these systems [51, 73], and conducted ethnographic studies around expectations and responses to these systems [47].

### 2.2 Governance of Content Moderation

Governance mechanisms regarding content moderation vary around the world [33]. For instance, in the United States, Section 230 of the Communications Decency Act of 1996 remains a crucial piece of legislation impacting content moderation. The legislation intended to reform tort law in light of the widespread impacts of the internet [69]. In a world becoming increasingly digitally connected, owners of digital platforms that connected massive amounts of internet users - either for purely social reasons or for commercial exchanges - were afraid that they might face an untenable burden of liability for the behaviors of their user base. In response to these fears, Section 230 effectively removes liability for these internet companies - called “intermediaries” - in all cases except those in which they “materially contribute” to this harmful user content.<sup>1</sup> While this doctrine allowed online media and commerce to flourish without the constant threat of costly law suits, it also fostered an environment that was, at best, neutral to those harmed by user generated content. By completely detaching these companies from the harms stemming from their users’ behaviors, a poor system of incentives developed in which genuinely harmful content proliferated.

Social media companies have, to date, used a largely tactical and supposedly neutral approach to content moderation, characterized by an ever-growing list of ad-hoc rules formed in response to political pressure from lawmakers or interest groups. This, according to Wihbey et al. [76], is unjust and “creates an uncertain epistemic environment ... inviting claims of bias, favoritism, and censorship. [p. 8]” As an example of the limited use and potential harms of such ad-hoc tactics, consider Facebook’s response to the claim, especially made by the politically conservative, that misinformation

labels displayed by social media companies are politically biased. In response to this, Facebook is reported to have changed some of its content standards to allow content that previously would have been flagged as misinformation to remain unlabeled and unaltered [23]. That is, the response consisted of altering content moderation standards *as they relate to only some sources of information* to mitigate the claim. Yet, this response could lead to potential discrepancies in the quality of labeled misinformation for different seekers of information, and the harms associated with this will be discussed in subsequent sections.

### 2.3 Harms of Misinformation

Many significant harms associated with misinformation and disinformation have been documented. For instance, researchers have observed harms related to: the adoption of addictive habits [77]; receiving misleading health advice during a pandemic [75]; democracy and social institutions [55]; beliefs about climate change [72]; and situational awareness in humanitarian crises [71].

Search engines and social media platforms may utilize targeted advertising as a source of revenue, in which content creators can promote their content to users with certain demographic or personality characteristics that yields a higher propensity to engage with the content [41, 43]. Research on “computational propaganda” has highlighted how adversaries intent on causing societal strife have used targeted advertising to show “weaponized” information likely to manipulate specific groups of users [14, 61]. (Some have also called this domain of research “social cyber-security” [12].) Importantly, in some cases such attacks are facilitated by the automation of key tasks by social media platforms. For instance, there is evidence that algorithms implemented by Facebook created categories that clustered users based on their propensity to engage with anti-Semitic articles, and this information was made accessible to those willing to pay for it [4].

In addition to being a potential risk to democracy and national security, targeted advertising can also pose an *epistemic* threat to communities. This can happen, for instance, when the technology facilitates the promotion of fraudulent products and conspiracy theories. Reports regarding ads on Facebook for a hat that could “protect [one’s] head from 5G cell-phone radiation” are an instance of this threat, particularly when they coincide with the increased popularity and circulation of the related conspiracy theory on social media [53]. Further investigation of such ads point to the role of target advertising, such as when Facebook determines that a user is interested in the product category “pseudoscience” [53].

Our work is crucially different from these lines of research in that we do not seek to study the harms associated to misinformation itself, but rather the injustices that particular stakeholders may encounter when deploying algorithmic tools to combat this phenomenon.

### 2.4 Algorithmic Misinformation Detection

To combat the scale of misinformation, large social media companies (like Facebook) have developed algorithms and policies to identify and act upon misleading information shared on their platform. Facebook’s system, in particular, automatically detects factual

<sup>1</sup>Jones v. Dirty World Entertainment Recordings LLC, 755 F.3d 398 (6th Cir. 2014)

claims worth checking. These claims would then be sent to an independent fact-checker to assess its veracity. Facebook also employs algorithms to determine if claims have been previously fact-checked by the human fact-checkers. If the result of this assessment is that a claim is false or misleading, Facebook “add(s) warnings and more context to content rated by third-party fact-checkers, reduc(es) their distribution, and remov(es) misinformation that may contribute to imminent harm” [26].

Algorithmic misinformation detection can be composed of many sub-tasks, which some systems tackle independently while others attempt to solve in an end-to-end fashion. While the specifics of these tasks may evolve and change over time, we draw from Guo et al. [34] to differentiate between three core (sequential) tasks: (1) *Check-worthiness*, which aims to spot factual claims that are worthy of fact-checking [11, 31, 39, 45], (2) *Evidence retrieval* of potential evidence for identified claims [21, 49, 56, 60, 66, 70, 74], and (3) *verdict prediction*, which aims to establish the veracity of a claim [60, 63, 74]. In a survey on the topic by Zhou and Zafarani [78], the authors identify how misinformation can be detected from four perspectives: (1) the false knowledge it carries; (2) its writing style; (3) its propagation patterns; and (4) the credibility of its source. In subsequent sections, we consider “general” machine learning methods with features that account for these four perspectives related to fact-checking.

### 3 JUSTICE & STAKEHOLDERS IN MISINFORMATION DETECTION

Misinformation detection pipelines consist of varied and many tasks, each of which can give rise to many ethical concerns. Precisely articulating these concerns requires a coherent normative framework for identifying the relevant stakeholders, mapping their legitimate rights and interests, and disentangling the distinct claims that they may justifiably have to justice. Therefore, well-known frameworks that focus on only a single stakeholder (i.e. the source of information) and single ethical notion (free speech) will be inadequate to describe the complex ethics of the information environment.

In this section, we provide this framework by building on Mathiesen’s [54] notion of *informational justice*—a multidimensional concept characterized in terms of “justice for persons and communities in their activities as seekers, sources, and subjects of information” [54, p. 199]. Mathiesen’s [54] framework originated in the context of information sciences, and so most directly applies to discussions of justice surrounding (*mis*)information. In order to comprehensively capture the concerns regarding misinformation *detection*, therefore, we need to appropriately extend the framework. In this section, we explain both the framework as well as our extensions. To see why a multidimensional notion of justice is needed in the first place, we begin by contrasting the situation in misinformation detection with the more familiar discussions of justice in allocation.

In common discussions of algorithmic fairness, each data point typically pertains to individuals (e.g., a job candidate in the case of hiring, a school district in the case of public funding allocations). These individuals tend to occupy a single direct stakeholder role,

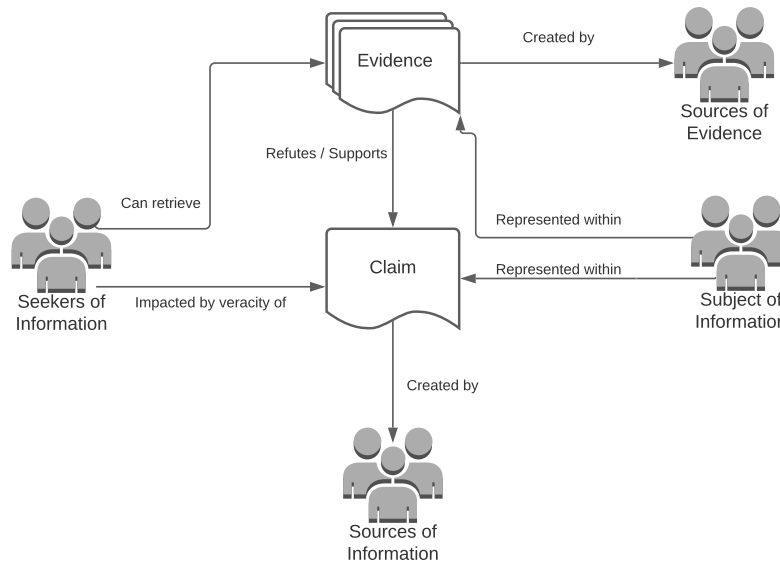
namely, that of decision subject. In allocation situations, these decision subjects encounter a decision-maker (e.g., an employer, a public funding body) who must decide how to distribute scarce resources and opportunities among them. Generally, here it is the decision subjects who have an invested interest in justice—more specifically distributive justice [24]—, while the primary concerns of the decision-makers (to whom the demands of justice apply) might be in values such as efficiency (e.g., for profit or for social welfare).<sup>2</sup>

In the case of informational items, in contrast, each data point typically relates to *multiple* stakeholders, who interact with or are directly implicated in the informational item. A tweet or a book, for example, has an *author* and an *audience*, and it can have a particular individual or group as its *subject matter*. In Mathiesen’s informational justice framework, these individuals or groups map onto three stakeholder roles: *seekers of information*, *sources of information*, and *subjects of information*. In the case of misinformation detection, in addition to the informational item that is the focus of evaluation (e.g., a tweet about which there has been a complaint), there is an additional type of informational item: *evidence* (e.g., in the form of the original complaint, or subsequent evidence gathered in verification). We identify an additional stakeholder role that is relevant to the evidence gathered during various stages of the process—from check-worthiness to truth verification): *sources of evidence*. Below, we describe each of these stakeholder types in more detail and explain the notion of justice most directly associated with them. Figure 1 maps out these stakeholders in relation to the two informational items involved in misinformation identification: a claim and evidence regarding the claim.

*Sources of Information and Participatory Justice.* Sources of information in misinformation detection are authors of articles or content creators that produced or else posted a claim for dissemination over social media. Conceptions of justice relating to the sources of information involve balancing the right to participate in the “marketplace of ideas” with the duty to act “appropriately” (often as dictated by either federated or centralized user agreements) and to not intentionally deceive others. The notion of justice relevant to the sources of information thus pertains to participation. *Participatory* justice entails that “all members of society should have opportunities to communicate their point of view alone or in concert with others, to have that point of view taken into account, and to take part in shared decision making about the provision of information resources.” In the context of misinformation detection, this entails certain rights of all users to have their opinions or thoughts heard by those connected to them, and, when applicable, to a broader audience. This would imply that all sources of information should be treated equally in a given social media ecosystem, regardless of their affiliation with any social-cultural group.

*Subjects of Information and Recognition Justice.* Subjects of information are represented within claims made by sources of information. Subjects can be individuals or groups, or they can be conceptual topics. When subjects are individuals or groups, a just system would ensure that their reputation faces no undeserved

<sup>2</sup>The decision-makers might also be interested in being just, and, when we go beyond direct stakeholders, particular communities and groups as well as general public might also prize distributive justice.



**Figure 1: This figure demonstrates how stakeholders are related to two informational items (claims and evidence) in automated misinformation detection, borrowing heavily from the informational justice framework posited by Mathiesen [54].**

harm from inaccurate claims made by others. More broadly, *recognitional justice* implies that “contents of information available within the information environment should include fair and accurate representations of all members of society.” In the current context, recognitional justice has particular importance, for example, in cases of false claims and derogatory language directed at particular social-cultural groups.

*Seekers of Information and Distributive Justice.* Seekers of information are people or groups that may be impacted by the presence of information about a topic that is of relevance to them and the veracity of the claims (tweet, article, images, etc) about that topic circulating online. Given the foundational ways it shapes and interacts with our resources and opportunities, information can be seen as a critical good in our societies. Accordingly, the notion of justice that is particularly pertinent in relation to seekers of information is largely *distributive* in nature. That is, justice here pertains to how informational resources are distributed among individuals and groups, and it is undermined by disparities in the availability or quality of relevant information. On social media, for example, such disparities can happen when users from certain demographic backgrounds do not see potentially relevant information (e.g., do not see a relevant job ad because of a clustering of users in homophilic networks) or receive lower quality information (e.g., receive ads about irrelevant jobs perhaps due bias in the recommendation system). Similar issues apply in the case of misinformation detection. Distributive injustice can happen, for example, when misinformation targeted at a marginalized community is more likely to circulate unhindered. Like other cases of distributive justice, what precisely the demands of justice consist in here requires specifying (and agreeing upon) the relevant resource that is being distributed (e.g.,

information, “quality” of information in some sense) as well as the appropriate rule for the distribution of that resource.<sup>3</sup>

*Sources of Evidence and Epistemic Injustice.* Individuals generating the second informational item—that is, evidence—constitute an additional stakeholder. Depending on the system design and the stage of the misinformation detection pipeline, this stakeholder may be a user of a platform who flags a claim (e.g., a twitter post), a fact-checker who independently analyzes the claim, a crowd worker surveyed to assess the veracity of the claim, or the author of an article that provides a stance towards the claim. In each case, the activities of these individuals results in the generation of a piece of evidence about the focal claim. Insofar as their activities directly shapes the information ecosystem, *sources of evidence* also may have a claim to participatory justice.

Here, however, we want to focus on another source of (in)justice that applies to sources of evidence more specifically, namely, *epistemic justice*. To the extent that misinformation detection systems fail to incorporate or else disregard the relevant knowledge produced by some sources of evidence, they may result in epistemic injustice. Particularly relevant in the misinformation detection context might be cases of *testimonial* injustice, where the claims of a source of evidence (e.g., that a tweet constitutes hate speech) receives an unfair deficit of credibility downstream (from other sources of evidence or an algorithm), and diminishes the source’s capacity as a *knower* [30]. This type of epistemic injustice can be particularly problematic if it systematically affects members of marginalized communities.

Epistemic justice is not the only type of justice that is relevant to sources of evidence. Depending on a source of evidence’s role, considerations of distributive justice can also come into play. Consider,

<sup>3</sup>The appropriateness of the rule might depend on the context.

for example, the emotional toll of hate speech on individuals from the targeted group. If there is a disparity in the amount and circulation of such speech such that individuals from affected groups are more likely to be forced to report and flag these items (that is, to voluntarily assume the role of sources of evidence), to that extent the information platform exhibits distributive injustice, potentially compounding existing disparities. Similar points apply to the toll of such speech on the mental health of crowdsourced workers who are tasked to label or else assess the veracity of the claim.

In this section, we primarily focused on notions of justice in relation to *claims* (or evidence pertaining to them) that form the input to misinformation detection systems. In addition to this, issues of distributive justice become salient in relation to the *misinformation detection systems* themselves, insofar as the performance and impacts of these systems shapes how the benefits (e.g. information quality control) and burdens (e.g. over-scrutiny) of the technology are distributed. For example, as highlighted in 2.1, in many instances the classification errors of existing algorithms disproportionately harm members of disadvantaged groups, thus resulting in concerns about distributive injustice.

## 4 SITUATING ALGORITHMIC INJUSTICE IN THE MISINFORMATION DETECTION PIPELINE

Machine learning can be employed to automate the entire misinformation detection pipeline, or to assist in specific stages of it. Figure 2 identifies the three central stages of the process, and provides a generalized view of the target objective at each of them. While specific design choices may vary across systems (e.g., choosing if check-worthiness is a multi-class label prediction problem or a regression problem), this general formulation enables us to situate the different types of injustices that may emerge at distinct functional stages of these systems.

Studying, anticipating or mitigating algorithmic harms requires the identification of stakeholder(s) affected by incorrect algorithmic predictions, together with the notion(s) of justice relevant to a stakeholder when a certain type of error is made. For example, some errors might result in undue scrutiny of speech or disproportionate harms stemming from the spread of misinformation, depending on whether the algorithm is over- or under-scrutinizing the content relevant to a particular group. In this section, we discuss the potential harms associated with different algorithms in automated misinformation detection, grounded on the taxonomy introduced in Section 3. For each algorithmic stage in the pipeline, we discuss: (1) a generalized framework/objective for the machine learning task; (2) the relevant sense of injustice to different stakeholders when affected by errors in algorithmic decisions, and (3) an overview of potential sources of these injustices.

### 4.1 Automated Claim Detection

**4.1.1 The Machine Learning Task.** The first task in the automated misinformation detection pipeline is identifying “check-worthy” claims. According to Hassan et al. [39], check-worthy claims are claims for which the “general public” would be interested in assessing their validity. For fact-checking organizations, these would be the claims (out of broader possible universe of claims) that they

choose to evaluate. Snopes, for instance, says this about how they decide what to fact-check: “We have long observed the principle that we write about whatever items the greatest number of readers are asking about or searching for at any given time, without any partisan considerations” [67]. It is worth noting that the motivations for identifying a subset of check-worthy claims may vary depending on the context and system design. Most commonly, check-worthiness is a prioritization mechanism to allocate the scarce resource that is human fact-checkers. Given the scalability of algorithmic fact-checking, such scarcity of fact-checking resources could be altered, but check-worthiness may remain as a relevant mechanism to screen content that should be fact-checked, conditioned on the nature of the claim.

Several model architectures have been proposed for detecting check-worthiness of textual data (e.g. [1, 11, 31, 39, 45]). Hassan et al. [39], for instance, gather numerous expert crowd-sourced labels for political speech sentences as either *Non-Factual Sentence* (opinion, belief), *Unimportant Factual Sentence* (facts that are not relevant to general public), and *Check-Worthy Factual Sentence* (facts that are relevant to general public). They then extract sentence sentiment, named entities, and tf-idf features from the vocabulary and optimize a support vector machine classifier predicting check-worthiness. As common in misinformation detection, this system relies on training data from a very specific, and likely not generalizable, corpus: political speeches.

CheckThat!, an annual misinformation prediction conference, has provided a dataset that contains 629 labeled English tweets pertaining to COVID-19 as a benchmark dataset [11]. The researchers gathered survey responses to five questions that attempted to capture the check-worthiness of each tweet. Participants in the machine learning competition were then asked to aggregate these labels and develop a model to predict check-worthiness for novel tweets.

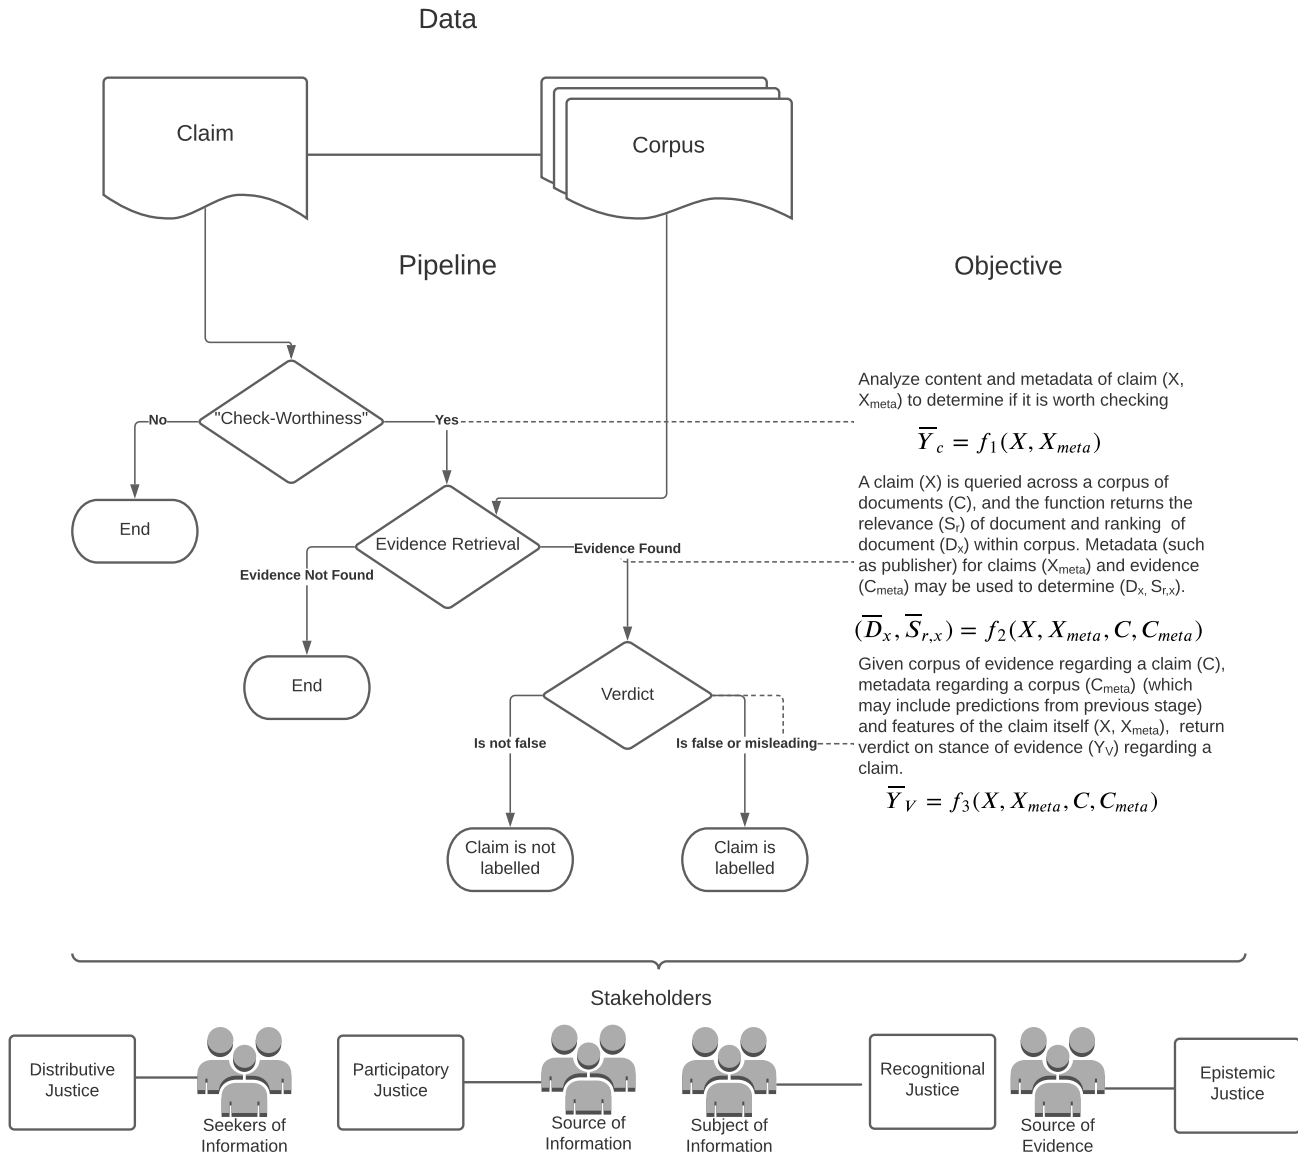
Researchers investigating COVID-19 misinformation adopted a very similar approach to Barrón-Cedeño et al. [11] for crowd-sourcing labels to detect claims made on social media, except instead of aggregating the label, they implemented a multi-task learning approach, in which the correlated survey questions used to determine check-worthiness were modelled jointly [1].

Some researchers have also framed the task as rumor detection, and have found success using metadata on social media (author information, rate of decay of engagement, reactions, and replies) jointly with semantic representations of the content to spot early propagation of rumors that should be fact-checked [31].

In general, we frame the machine learning task as trying to recover the function:

$$\bar{Y}_c = f_1(X, X_{meta}) \quad (1)$$

Where  $\bar{Y}_c$  is the predicted check-worthiness label, score or vector of labels (binary, multi-class, or a tuple in the context of multi-task learning) for a given claim,  $X$  is a semantic representation of the claim, and  $X_{meta}$  is a vector of metadata associated with the claim (likes, reposts, etc). If a piece of content is predicted to be check-worthy, then it will pass to further stages in the misinformation detection pipeline, where evidence will be gathered in order to produce a veracity score or label. The check-worthiness algorithm



**Figure 2: A generic framework, used in conjunction with Figure 1, to examine issues of justice regarding algorithmic decision-making in the automated misinformation detection pipeline.**

is concerned with reducing the scale of information for which evidence retrieval is necessary.

**4.1.2 Issues of Justice.** Under-scrutiny of the claims reported by, of interest to, or about a particular group can lead to a distributive injustice for sources of evidence, seekers and subjects of information, as a key moderator of the informational good (fact-checking) is not being distributed evenly amongst groups. Seekers of information may be harmed because they receive a comparatively lower standard of information quality, allowing them to more directly

feel the harms of misinformation than others. Recall the example from Section 2, where Facebook’s seemingly ad-hoc content moderation policy led to claims of bias from conservative sources of information—a claim of participatory injustice in our framework. In response, Facebook changed the standards for evaluating misinformation from conservative sources of information, reducing scrutiny of some of these sources. By changing the standard for misinformation detection for specific groups (in this case, conservatives), harms fall upon (conservative) seekers of information, who are now

more likely to be exposed to misinformation. Under-scrutiny errors at the claim detection stage may also affect subjects of information, who could be harmed if false claims were made against a community or an individual that could damage their reputation or fuel stigma and stereotypes and these claims were not prioritized to be fact-checked. Thus, failing to assess these claims could result in a recognitional injustice, by enabling inaccurate representations of members of society. Finally, in settings where user reports are an input to check-worthiness assessment, a disproportionate lack of attention to claims flagged by a community could constitute an epistemic injustice, which may additionally result in some of the injustices listed above for the same individuals and/or communities, when their role in relation to a claim is multi-faceted. While the type of injustice stemming from under-scrutiny will be different for different stakeholders, they all correspond to false negative errors. Thus, assessing *equal opportunity* [38] as an algorithmic fairness metric is of relevance, where the sensitive attribute category will vary conditioned on the stakeholder.

Over-scrutiny may also constitute harms, although these will depend on the interaction with other portions of the system. By itself, marking something to be verified does not necessarily harm sources, seekers or subjects, especially if we assume perfect accuracy in the fact-checking stage. However, that is an unrealistic assumption, and it is necessary to consider fairness under composition [20]. In particular, it has been shown in the algorithmic fairness literature that making the same type of errors, e.g. false negatives, when the output of one algorithm determines the input of the following one will result in compounding imbalances, which may correspond to compounding injustices [18]. Thus, monitoring *false positive rates* associated with predicted check-worthiness is especially pressing in conjunction with the errors at other stages of the pipeline.

**4.1.3 Sources of Harms.** Check-worthiness is an ambiguous task, and those designing a system must first decide how to define what being check-worthy means, and establish a quantitative approximation to it. This problem formulation stage may already induce bias. For example, check-worthiness may be defined as what interests the “general public” [39], or what “the greatest number of readers are asking about” [67]. Clearly, a simple focus on majorities disregards minority communities along multiple axes. Misinformation targeting specific minorities, e.g. an immigrant community, will by definition not be considered check-worthy unless it also affects other groups, constituting a distributive injustice for this group of seekers of information. Similarly, when these communities participate in the ecosystem as sources of evidence, by reporting or searching for certain items, their input will be neglected unless it matches the interest of majority groups, resulting in an epistemic injustice.

Thus, the harms associated with false negative errors, or under-scrutiny of claims, may be disproportionately concentrated in certain groups as a result of *measurement bias* in the training labels [42]. Measurement bias—when an inappropriate proxy is used to measure a construct [68]—may result from a choice of metric, e.g. the most frequently searched items, or from ambiguity in the formulation of survey questions used to collect check-worthiness labels. For instance, questions inquiring whether something is of interest to

the “general public” may elicit very subjective responses regarding who the “general public” is, and may downplay claims that are highly relevant (with potentially high stakes) to a particular minority group. This issue may be exacerbated for subjects of information if there is a cultural disconnect between those evaluating the claim (the sources of evidence) and the subject of information. If there is a cultural disconnect, those evaluating the content may fail to appropriately understand it as well as its implications or implied sentiment, as these may all vary contextually across cultures. For example, social stereotypes affect how people understand language directed toward minority groups, leading to higher false negative assessments when labelers have negative stereotypes about the subject of information [16]. One direction for future research would be to evaluate the impact of cultural alignment between the subject of (mis)information and the source of evidence to see if there are significant differences between a culturally-aligned and culturally-misaligned subject-source pairs. If one were to find differences between the subject-source pairs, this might imply that culturally-misaligned annotators provide lower signal (or high bias) labels.

## 4.2 Evidence Retrieval

**4.2.1 Machine Learning Task.** The next step in the pipeline involves retrieving evidence to either support or reject the claims that are deemed “check-worthy”. Two main types of approaches have been proposed for this task. First, some methods (such as FAKTA [60]) utilize Google’s commercial API to automatically search claims and retrieve (potentially) relevant evidence from the broader internet corpus. In some instances, this is followed by a post-processing steps in which results retrieved from search engines are merged with a database measuring the credibility of sources of information, and only evidence from sources deemed credible are kept [60]. The second type of approach is especially tailored to curb the spread of misinformation that has been previously fact-checked by humans. To do so, it uses a corpus of previously fact-checked claims stemming from either crowd-sourced knowledge, such as Wikipedia [70] or surveys [2], or from professional assessments, such as Snopes or PolitiFact [66]. The task is then to map claims to previously fact-checked statements.

In general, we can say that given a claim  $X$  and a corpus  $C$ , each of which may have associated metadata ( $X_{meta}$  and  $C_{meta}$ ), the task is to retrieve a ranked list of the documents in the corpus  $\bar{D}_x$ , each with an associated relevance score  $\bar{S}_x$ , as shown in Equation 2.

$$(\bar{D}_x, \bar{S}_x) = f_2(X, X_{meta}, C, C_{meta}) \quad (2)$$

**4.2.2 Issues of Justice.** Typically, information retrieval algorithms are useful if they have both high recall and precision, as these are indicators of relevance and completeness of retrieved information, respectively [52]. If returned evidence is disproportionately incomplete or irrelevant for some groups, then issues of justice may arise. Below, we posit several ways stakeholders might face issues of justice in automated evidence retrieval.

As previously noted, since the broader internet contains many documents making claims that have not been verified themselves, system designers may consider implementing an intermediate step to improve reliability of returned evidence: constraining the corpus

of potential evidence to “trusted” sources, such as Wikipedia, PolitiFact, or other verified bodies of knowledge. However, these trusted sources of evidence may not engage with or have knowledge of claims from all peoples, and could exhibit bias in which claims they have chosen to fact-check. This could result in irrelevant and/or incomplete evidence gathered at this stage in the pipeline. Measures of recall could thus be relevant for assessing harms to seekers of information and sources of evidence. Seekers of information could be harmed if the corpus is limited to trusted sources of evidence that do not attend to some groups and communities, yielding low recall of information relevant to them. For sources of evidence, this may constitute an epistemic injustice, as failing to include them in the corpus of “credible sources” unjustly diminishes its value and disregards the knowledge they have produced.

Harms may also stem from the use of commercial search engines to retrieve evidence (with or without the intermediate step previously mentioned). Commercial search engines may down-weight evidence from websites with little traffic or poor search engine optimization (SEO), yielding different rankings for sources with the same relevance  $S_x$  as a result of maximizing for another commercial metric (click-through rate, time spent per link). If websites containing relevant information to support or reject claims from marginalized communities appear lower in search results despite having the same relevancy, this may inadvertently reduce the relevancy or completeness for the returned corpus of evidence ( $C$ ), which would again be reflected in recall metrics. This will constitute an epistemic injustice against sources of evidence by dismissing or discounting the information they produce in a way that does not hold an appropriate relationship to the quality of the information. Further, these issues may be exacerbated if only a relatively small corpus of potential evidence is kept after the claim is queried (e.g. first page of results) before document ranking and relevancy to the claim are determined through other means. Conversely, if websites that promote misleading evidence are ranked higher due to higher predicted engagement (i.e. “click-bait”) rather than their relevance to the claim, this could render low precision in the evidence retrieval system. Such low precision metrics in relation to different types of information could potentially harm seekers and subjects of information; seekers could be affected through degraded fact-checking quality, while subjects could suffer recognitional injustices if false statements about them are encoded as part of a corpus treated as veridical.

**4.2.3 Sources of Harms.** Concerns over disparities in evidence retrieval are primarily grounded in properties of the current fact-checking ecosystem [48] and dynamics of the way information is gathered and disseminated on the web [7]. These may bias the composition of the corpus  $C$ , metadata regarding the corpus  $C_{meta}$ , or the proxies used to estimate  $(\bar{D}_x, \bar{S}_x)$ .

While professional fact-checkers tend to deliver gold-standard evidence, they are significantly concentrated in North America and Europe, especially when accounting for population [48], as shown in Table 1. Therefore, due to *historical bias* in the data-generating process, there may be less previously fact-checked evidence available on the web for claims relevant to people from Asia, South America, and Africa. If claims regarding non-Western subjects of

information *are* fact-checked by a Western fact-checking organization, it will be examined from the perspective of Western experience and with a Western analytical framework [37]; this may induce a *measurement bias* related to the stance of the fact-checking organization, and the contextual information they may or may not have access to. For instance, consider the stance a highly reputable news agency, *The New York Times*, took towards the claim that Iraq possessed weapons of mass destruction. The newspaper famously published information stating that Iraq was in possession of weapons of mass destruction, only to, some time later, acknowledge their lack of integrity when interviewing sources of information. They said that their reporting was “eagerly confirmed by United States officials convinced of the need to intervene in Iraq.” [22]

Another highly regarded source of evidence that is widely incorporated into fact-checking systems is Wikipedia [70]. While the premise of Wikipedia as an open-source encyclopedic body may yield hopes of a lower bias environment that is collectively constructed, in an early version of English Wikipedia, only 0.04% of authors were responsible for the development (i.e. creation and related research) of over half the body of knowledge [7]. Thus, “the notion that it represents the wisdom of the overall crowd is an illusion” [p. 56]. This effect has been called an *activity bias*, and can impact the objectivity of open-source evidence if relevant groups do not contribute to the production of evidence. Furthermore, the amount of information available varies by subject and language. For instance, in a study comparing Polish and American Wikipedia articles about famous people, researchers noted that, “English language entries have more references and external links... They also tend to be longer than Polish language entries” [15].

When search engines are used to retrieve information, *learning bias* affecting the information retrieval algorithms may in turn bias the evidence retrieval system in undesirable ways. For example, if the commercial search engine used to retrieve the documents has the goal of optimizing revenue, using this as a ranking of relevance could have detrimental effects that could disproportionately affect sources of evidence, as well as seekers and subjects of information.

*Measurement bias* may also emerge when relying on crowd-sourced indicators of credibility. For example, FAKTA [60] incorporates metadata ( $C_{meta}$ ) regarding the credibility of sources of evidence, which is ultimately used to algorithmically weight the stance of evidence towards a claim. The metadata is pulled from an existing, crowd-sourced database<sup>4</sup> that contains ratings on roughly 4,300 sources of evidence. If credibility is determined by a majority group, sources of evidence that are reliable but unknown or mis-assessed by the majority group will have their evidence down-weighted, even if the evidence is truthful and relevant to the claim. Future work should assess whether databases like this contain cultural biases in their assessment and composition of sources of evidence, which could bias  $(\bar{D}_x, \bar{S}_x)$ .

## 4.3 Automated Verdict Generation

**4.3.1 Machine Learning Task.** After a claim has been determined to be “check-worthy” and relevant evidence has been gathered about the claim, the final phase involves generating a verdict regarding the claim. This is the most visible and consequential phase in the

<sup>4</sup>[www.mediabiasfactcheck.com](http://www.mediabiasfactcheck.com)



Continent	Fact-Checking Organizations	Population (Billions)	Fact-Checking Organizations Per Billion
Africa	34	1.37	24.76
Asia	87	4.68	18.59
Australia	5	0.04	119.05
Europe	99	0.75	132.53
North America	81	0.60	135.91
South America	39	0.43	89.86

**Table 1: This table shows the number of professional fact-checking organizations per billion population across 6 continents. The concentration of fact-checking organizations is roughly 7.3x higher in North America than in Asia when population is taken into account. (Source: [48])**

pipeline, as there is likely to be some significant action taken by the platform if a verdict is reached that the claim is false or misleading. Such systems may aim to provide labels visible to seekers of information, to assist with content moderation, or to provide decision support to human fact-checkers.

An important subtask that some have found relevant to generating a verdict is predicting the stance of each piece of relevant evidence towards the claim. In the standard version of this task, a document from the corpus  $C$  is probabilistically labeled (e.g. Supports, Against, Neither), depending on how it relates to the semantic representation of the claim  $X$ . For more information on this task, see Küçük and Can [46]. The predicted stance can then be viewed as metadata in relation to each evidence document, serving as input to a verdict generation algorithm. For example, FAKTA’s end-to-end automated fact-detection system uses the stance and credibility of supporting pieces of evidence to yield a final prediction for the veracity of a claim [60].

Other approaches have considered estimating the verdict of a claim without the need to retrieve supporting evidence. For example, Reis et al. [63] construct semantic features extracted from the claim meant to capture subjectivity of language, features about the publisher of the claim meant to capture possible bias, as well as social media metadata surrounding the claim such as the number of likes and sharing patterns.

In a generalization of Vo and Lee [74], the machine learning task can be formulated as follows:

$$\bar{Y}_v = f_3(X, X_{meta}, C, C_{meta}) \quad (3)$$

In this equation, the verdict  $\bar{Y}_v$  is functionally related to the semantic representation of claim  $X$ , metadata of the claim  $X_{meta}$ , a corpus of evidence  $C$ , and the metadata regarding the corpus  $C_{meta}$ , such as credibility and the predicted rank, relevance, and stance of each piece of evidence in relation to the claim. Naturally and as in the previous formulations, not all systems need to consider all inputs. The training data labels  $\bar{Y}_v$  could be collected from professional fact-checkers, such as Snopes or PolitiFact, or they could be crowd-sourced as well, as in Babaei et al. [5] and Allen et al. [2]. If there is more than one observation per claim, the labels from multiple crowd-sourced or professional fact-checkers are aggregated in some way (e.g. the mean, median, or ideological difference) to create a composite “ground truth” label [5].

**4.3.2 Issues of Justice.** Injustices to stakeholders occur when the predicted label  $\widehat{Y}_v$  conflicts with the true, and potentially unobserved, label  $\bar{Y}_v$ . In the case of a *false negative*, in which no action (like labeling) is performed on potentially misleading information, significant harms are likely to be felt by seekers and subjects of information. As mentioned in Section 2, harms to seekers of misinformation are vast, and constitute broad threats to one’s health, happiness, and societal institutions, especially if the misinformation is adversarial in nature (i.e. disinformation). Therefore, groups of seekers exposed to low quality information due to incorrect predictions face a distributive injustice of high quality of content moderation. Subjects of information, likewise, suffer recognitional injustice if their reputation is harmed by damaging and misleading information shared without a warning label, as seekers may interpret the claim to be true. This harm can be especially great to individual subjects of information in the court of public opinion. For instance, subjects of rumors spread on WhatsApp had their reputation damaged to such an extent that locals incorrectly thought they posed a grave threat, formed a mob, and murdered them [64].

A *false positive* occurs when a moderation action is incorrectly taken against a claim made on a platform. When false positives are concentrated within some group of sources of information conditioned upon a particular sensitive attribute, this constitutes a participatory injustice. This is because the sources cannot contribute to or participate in the broader informational exchange at the same rate as others. If the action taken against the claim is to place a warning label on it, research has shown that this will reduce the credibility of the content [44]. Therefore, their credibility and capacity as a participant in the informational exchange has been unjustly diminished. Additionally, we can assume that a high rate of false positives will yield some reputational harm to the source of information, as their reputation might be damaged amongst their peers for sharing content that is routinely labeled as misleading. This harm would only be magnified if the source of information is also represented either expressly or contextually as the subject of information. For example, in the context of hate speech detection, a higher rate of false positives for posts in which a source of information discusses their identity of belonging to a disadvantaged group [19] affects stakeholders in their roles as both sources and subjects. Additionally, similar in nature to the harms associated with false negatives, seekers of information may be damaged when an action is incorrectly taken against a claim, especially if the action

in consideration involves removing the claim from the platform. Seekers have a right, fundamentally, to view legitimate information unhindered, and this could qualify as a distributive injustice to access to legitimate information.

**4.3.3 Sources of Harms.** Research related to biases in determining the veracity of claims has covered two main topics, both related to the “wisdom of crowds” in determining misinformation from crowdsourcing labels. It is worth noting that existing research in this space has only assessed bias in the labels themselves, often considering crowdworkers as potential fact-checkers rather than as sources of training data for algorithms. Ways in which these biases affect algorithmic predictions are yet to be studied in this domain.

One question that has been examined is whether obtaining labels from laypeople induces a *measurement bias* when determining what constitutes misinformation. Allen et al. [2] find that, while individual laypeople do tend to show considerably more bias and inconsistency than experts, a relatively small, politically balanced crowd yields an average label that is highly correlated to the average label of a few professional fact-checkers. This seems to imply that, under the stated circumstances, the measurement bias in the training labels is mitigated if the set of labelers is balanced with respect to a sensitive attribute  $A = \textit{political affiliation}$ . This finding suggests that there is wisdom in crowds, and provides a criteria for selecting the composition of layperson labeling teams. However, this research focuses largely on American political claims, so it is unclear whether a crowd of politically balanced laypeople maintain unbiased labels when the subject of information is non-Western or relatively unknown.

Another important question to address is whether *aggregation bias* exists when crowdsourced labels are averaged. In other words, does the crowd provide more signal beyond an “average perception of truth”? By comparing crowdsourced perceptions to labels generated from professional fact-checking websites (such as Snopes), researchers have proposed several new signal-rich measures of “perceptual difference” amongst the crowd [5]. For instance, they look at the marginal distributions of crowdsourced truth perceptions by political affiliation, which they refer to as “ideological mean perception bias” or IMPB. This is a promising step in assessing how differences between people affect their perceptions of truth in news. However, there may be differences in perceptions that extend beyond political ideology. Crucial to our study, this metric may be a useful way to compare perceptions of truth from groups of any protected attribute (race, gender, ethnicity, etc) to determine if there is a high level of disagreement between members of a particular group and professional fact checkers. If there are significant differences, the next logical step is to ask why this is the case. Does this group commonly encounter misinformation that has altered their perception? Or does this group have access to unique knowledge that gives them more expertise on a given topic than a general fact-checker? Some labelers may have epistemically advantaged standpoints [29], and aggregation mechanisms such as majority voting may fail to represent their assessments [17].

## 5 CONCLUSION

In this paper, we employed and extended the informational justice framework to appropriately consider justice in algorithmic misinformation detection from the perspective of four key stakeholders: *seekers of information*, *sources of information*, *subjects of information*, and *sources of evidence*. Grounded in notions of *distributive*, *participatory*, *recognition*, and *epistemic* justice, we analyzed the harms that might impact various stakeholders if algorithmic errors fall disproportionately on any group with a shared sensitive attribute. Specifically, we conducted this analysis across the three (generalized) stages of the automated pipeline: (1) *check-worthiness*; (2) *evidence retrieval*; and (3) *verdict prediction*. We also put forth concrete ways to measure these harms in automated systems by drawing connections to widely used measures of algorithmic bias, and dissected potential sources of these biases.

In this work, the focus of our analysis has been the algorithmic prediction stages. However, downstream effects and dynamic human-algorithm interactions may yield novel concerns. For instance, a growing body of work has analyzed the effects of different interventions meant to curb misinformation [6, 35, 59]. Thus, considering algorithmic errors in conjunction with the psychological effects of different interventions may be of importance, especially as it concerns the design and deployment of explainable fact-checking systems.

A crucial challenge for algorithmic auditing of misinformation detection algorithms are the difficulties defining and labeling sensitive groups for the different stakeholders. In some cases, sources of information and/or evidence may have clear identities of belonging, e.g. when considering news media, Univision can be identified as Latinx media and Al Jazeera as a Middle Eastern source. Similarly, subjects of information can often be identified by leveraging computational linguistics to identify the subject of a claim when a group is explicitly mentioned in the claim. However, in some cases this is not straightforward and may have important ethical concerns associated with it [32]. For example, attempting to infer gender can result in harms to transgender individuals [36].

In information ecosystems, justice is a multidimensional notion. The contributions presented in this paper provide a coherent framework for designers, policy makers and researchers to precisely articulate and identify the varied ways in which concerns about injustice can arise in algorithmic misinformation detection. In doing so, it offers guidance to efforts towards developing technical responses that counter the threat of misinformation in an equitable manner.

## FUNDING ACKNOWLEDGMENT

This work was supported in part by Good Systems, a research grand challenge at the University of Texas at Austin, and by a Google AI Award for Inclusion Research.

## REFERENCES

- [1] Firoj Alam, Shaden Shaar, Alex Nikolov, Hamdy Mubarak, Giovanni Da San Martino, Ahmed Abdelali, Fahim Dalvi, Nadir Durrani, Hassan Sajjad, Kareem Darwish, and Preslav Nakov. 2020. Fighting the COVID-19 Infodemic: Modeling the Perspective of Journalists, Fact-Checkers, Social Media Platforms, Policy Makers, and the Society. *CoRR abs/2005.00033* (2020). arXiv:2005.00033 <https://arxiv.org/abs/2005.00033>

- [2] Jennifer N L Allen, Antonio A Arechar, Gordon Pennycook, and David G Rand. 2020. Scaling Up Fact-Checking Using the Wisdom of Crowds. (Oct 2020). <https://doi.org/10.31234/osf.io/9qdzta>
- [3] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine bias. *ProPublica*, May 23, 2016 (2016), 139–159.
- [4] Julia Angwin and Madeleine Varner. 2017. Facebook enabled advertisers to reach 'Jew Haters'. *ProPublica* (Sept 2017). <https://www.propublica.org/article/facebook-enabled-advertisers-to-reach-jew-haters>
- [5] Mahmoudreza Babaei, Juih Kulshrestha, Abhijnan Chakraborty, Elissa M. Redmiles, Meeyoung Cha, and Krishna P. Gummadi. 2021. Analyzing Biases in Perception of Truth in News Stories and Their Implications for Fact Checking. *IEEE Transactions on Computational Social Systems* (2021), 1–12. <https://doi.org/10.1109/TCSS.2021.3096038>
- [6] Sumitra Badrinathan. 2021. Educative interventions to combat misinformation: Evidence from a field experiment in India. *American Political Science Review* 115, 4 (2021), 1325–1341.
- [7] Ricardo Baeza-Yates. 2018. Bias on the Web. *Commun. ACM* 61, 6 (may 2018), 54–61. <https://doi.org/10.1145/3209581>
- [8] Michele Banko, Brendon MacKeen, and Laurie Ray. 2020. A unified taxonomy of harmful content. In *Proceedings of the fourth workshop on online abuse and harms*. 125–137.
- [9] Solon Barocas, Kate Crawford, Aaron Shapiro, and Hanna Wallach. 2017. The problem with bias: Allocative versus representational harms in machine learning. In *9th Annual Conference of the Special Interest Group for Computing, Information and Society*.
- [10] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [11] Alberto Barrón-Cedeño, Tamer Elsayed, Preslav Nakov, Giovanni Da San Martino, Maram Hasanain, Reem Suwaileh, Fatima Haouari, Nikolay Babulkov, Bayan Hamdan, Alex Nikolov, Shaden Shaar, and Zien Sheikh Ali. 2020. Overview of CheckThat 2020: Automatic Identification and Verification of Claims in Social Media. *CoRR abs/2007.07997* (2020). arXiv:2007.07997 <https://arxiv.org/abs/2007.07997>
- [12] David M Beskow and Kathleen M Carley. 2019. *Social cybersecurity: an emerging national security requirement*. Technical Report. Carnegie Mellon University Pittsburgh United States.
- [13] Su Lin Blodgett, Solon Barocas, Hal Daumé III, and Hanna Wallach. 2020. Language (technology) is power: A critical survey of "bias" in nlp. *arXiv preprint arXiv:2005.14050* (2020).
- [14] Gillian Bolsover and Philip Howard. 2017. Computational Propaganda and Political Big Data: Moving Toward a More Critical Research Agenda. *Big Data* 5, 4 (2017), 273–276. <https://doi.org/10.1089/big.2017.29024.cpr> PMID: 29235917
- [15] Ewa S. Callahan and Susan C. Herring. 2011. Cultural bias in Wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology* 62, 10 (2011), 1899–1915. <https://doi.org/10.1002/asi.21577> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.21577
- [16] Aida Mostafazadeh Davani, Mohammad Atari, Brendan Kennedy, and Morteza Dehghani. 2021. Hate speech classifiers learn human-like social stereotypes. *arXiv preprint arXiv:2110.14839* (2021).
- [17] Aida Mostafazadeh Davani, Mark Diaz, and Vinodkumar Prabhakaran. 2021. Dealing with disagreements: Looking beyond the majority vote in subjective annotations. *arXiv preprint arXiv:2110.05719* (2021).
- [18] Maria De-Arteaga, Alexey Romanov, Hanna Wallach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnamurthy Kenchadapadi, and Adam Tauman Kalai. 2019. Bias in bias: A case study of semantic representation bias in a high-stakes setting. In *proceedings of the Conference on Fairness, Accountability, and Transparency*. 120–128.
- [19] Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. Measuring and mitigating unintended bias in text classification. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. 67–73.
- [20] Cynthia Dwork and Christina Ilvento. 2018. Fairness under composition. *arXiv preprint arXiv:1806.06122* (2018).
- [21] Nima Ebadi, Mohsen Jozani, Kim-Kwang Raymond Choo, and Paul Rad. 2021. A Memory Network Information Retrieval Model for Identification of News Misinformation. *IEEE Transactions on Big Data* (2021), 1–1. <https://doi.org/10.1109/TBDATA.2020.3048961>
- [22] NY Times Editors. 2004. FROM THE EDITORS; The Times and Iraq. *New York Times* (May 2004). <https://www.nytimes.com/2004/05/26/world/from-the-editors-the-times-and-iraq.html>
- [23] Sarah Ellison and Elahe Izadi. 2021. 'Definitely not the results we want': Facebook staff lamented 'perverse incentives' for Media. *The Washington Post* (Oct 2021). <https://www.washingtonpost.com/business/2021/10/26/conservative-media-misinformation-facebook/>
- [24] Jon Elster. 1992. *Local justice: How institutions allocate scarce goods and necessary burdens*. Russell Sage Foundation.
- [25] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press.
- [26] Facebook. 2020. Here's how we're using AI to help detect misinformation. <https://ai.facebook.com/blog/heres-how-were-using-ai-to-help-detect-misinformation/>
- [27] Facebook. 2021. Community Standards Enforcement. *Transparency Center* (Jun 2021). <https://transparency.fb.com/data/community-standards-enforcement/fake-accounts/facebook/#content-actioned>
- [28] Sina Fazelpour and David Danks. 2021. Algorithmic bias: Senses, sources, solutions. *Philosophy Compass* 16, 8 (2021), e12760.
- [29] Sina Fazelpour and Maria De-Arteaga. 2021. Diversity in Sociotechnical Machine Learning Systems. *arXiv preprint arXiv:2107.09163* (2021).
- [30] Miranda Fricker. 2007. *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- [31] Jie Gao, Sooji Han, Xingyi Song, and Fabio Ciravegna. 2020. RP-DNN: A Tweet level propagation context based deep neural networks for early rumor detection in Social Media. arXiv:2002.12683 [cs.SI]
- [32] Avijit Ghosh, Ritam Dutt, and Christo Wilson. 2021. When Fair Ranking Meets Uncertain Inference. *arXiv preprint arXiv:2105.02091* (2021).
- [33] Tarleton Gillespie. 2021. *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
- [34] Zhijiang Guo, Michael Sejr Schlichtkrull, and Andreas Vlachos. 2021. A Survey on Automated Fact-Checking. *CoRR abs/2108.11896* (2021). arXiv:2108.11896 <https://arxiv.org/abs/2108.11896>
- [35] Michael Hameleers. 2022. Separating truth from lies: comparing the effects of news media literacy interventions and fact-checkers in response to political misinformation in the US and Netherlands. *Information, Communication & Society* 25, 1 (2022), 110–126. <https://doi.org/10.1080/1369118X.2020.1764603> arXiv:https://doi.org/10.1080/1369118X.2020.1764603
- [36] Foad Hamidi, Morgan Klaus Scheuerman, and Stacy M. Branham. 2018. Gender recognition or gender reductionism? The social implications of embedded gender recognition systems. In *CHI*.
- [37] Thomas Hanitzsch. 2019. Journalism studies still needs to fix Western bias. *Journalism* 20, 1 (2019), 214–217. <https://doi.org/10.1177/1464884918807353> arXiv:https://doi.org/10.1177/1464884918807353
- [38] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016), 3315–3323.
- [39] Naeemul Hassan, Fatma Arslan, Chengkai Li, and Mark Tremayne. 2017. Toward Automated Fact-Checking: Detecting Check-Worthy Factual Claims by Claim-Buster (KDD '17). Association for Computing Machinery, New York, NY, USA, 1803–1812. <https://doi.org/10.1145/3097983.3098131>
- [40] Deborah Hellman. 2021. Big Data and Compounding Injustice. *Journal of Moral Philosophy, forthcoming, Virginia Public Law and Legal Theory Research Paper* 2021-27 (2021).
- [41] Ganesh Iyer, David Soberman, and J. Miguel Villas-Boas. 2005. The Targeting of Advertising. *Marketing Science* 24, 3 (2005), 461–476. <https://doi.org/10.1287/mksc.1050.0117> arXiv:https://doi.org/10.1287/mksc.1050.0117
- [42] Abigail Z Jacobs and Hanna Wallach. 2021. Measurement and fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 375–385.
- [43] Johannes Knoll. 2016. Advertising in social media: a review of empirical evidence. *International Journal of Advertising* 35, 2 (2016), 266–300. <https://doi.org/10.1080/02650487.2015.1021898> arXiv:https://doi.org/10.1080/02650487.2015.1021898
- [44] Timo Koch, Lena Frischlich, and Eva Lermer. 2021. The Effects of Warning Labels and Social Endorsement Cues on Credibility Perceptions of and Engagement Intentions with Fake News. (Jun 2021). <https://doi.org/10.31234/osf.io/fw3zq>
- [45] Lev Konstantinovskiy, Oliver Price, Mevan Babakar, and Arkaitz Zubiaga. 2020. Towards Automated Factchecking: Developing an Annotation Schema and Benchmark for Consistent Automated Claim Detection. (2020). arXiv:1809.08193 [cs.CL]
- [46] Dilek Küçük and Fazli Can. 2020. Stance Detection: A Survey. *ACM Comput. Surv.* 53, 1, Article 12 (feb 2020), 37 pages. <https://doi.org/10.1145/3369026>
- [47] Salla-Maaria Laaksonen, Jesse Haapoja, Teemu Kinnunen, Matti Nelimarkka, and Reeta Pöyhtäri. 2020. The datafication of hate: expectations and challenges in automated hate speech monitoring. *Frontiers in big Data* 3 (2020), 3.
- [48] Duke Reporters' Lab. 2016. Fact-checking Organizations. (Oct 2016). <https://reporterslab.org/fact-checking/#>
- [49] Nayeon Lee, Chien-Sheng Wu, and Pascale Fung. 2018. Improving Large-Scale Fact-Checking using Decomposable Attention Models and Lexical Tagging. (Oct-Nov. 2018), 1133–1138. <https://doi.org/10.18653/v1/D18-1143>
- [50] Rachel Lerman and Heather Kelly. 2020. Facebook says it labeled 180 million debunked posts ahead of the election. *The Washington Post* (Nov 2020). <https://www.washingtonpost.com/technology/2020/11/19/facebook-election-warning-labels/>
- [51] Kosisochukwu Madukwe, Xiaoying Gao, and Bing Xue. 2020. In data we trust: A critical analysis of hate speech detection datasets. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*. 150–161.
- [52] Christopher Manning. 2008. *Introduction to information retrieval*. Cambridge University Press, New York.

- [53] Nathalie Maréchal, Rebecca MacKinnon, and Jessica Dheere. 2020. *Getting to the Source of Infodemics: It's the Business Model: A Report from Ranking Digital Rights*. Technical Report. New America. 13–21 pages. <http://www.jstor.org/stable/resrep25417.5>
- [54] Kay Mathiesen. 2015. Informational Justice: A Conceptual Framework for Social Justice in Library and Information Services. *Library Trends* 64, 2 (2015), 198–225.
- [55] Spencer McKay and Chris Tenove. 2021. Disinformation as a Threat to Deliberative Democracy. *Political Research Quarterly* 74, 3 (2021), 703–717. <https://doi.org/10.1177/1065912920938143> arXiv:<https://doi.org/10.1177/1065912920938143>
- [56] P. Meesad. 2021. Thai Fake News Detection Based on Information Retrieval, Natural Language Processing and Machine Learning. *SN COMPUT. SCI* 2, 425 (2021). <https://doi.org/10.1007/s42979-021-00775-6>
- [57] Margaret Mitchell, Dylan Baker, Nyalleng Moorosi, Emily Denton, Ben Hutchinson, Alex Hanna, Timnit Gebru, and Jamie Morgenstern. 2020. Diversity and inclusion metrics in subset selection. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 117–123.
- [58] Shira Mitchell, Eric Potash, Solon Barocas, Alexander D'Amour, and Kristian Lum. 2021. Algorithmic fairness: Choices, assumptions, and definitions. *Annual Review of Statistics and Its Application* 8 (2021), 141–163.
- [59] 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).
- [60] Moin Nadeem, Wei Fang, Brian Xu, Mitra Mohtarami, and James R. Glass. 2019. FAKTA: An Automatic End-to-End Fact Checking System. *CoRR* abs/1906.04164 (2019). arXiv:1906.04164 <http://arxiv.org/abs/1906.04164>
- [61] Anthony Nadler, Matthew Crain, and Joan Donovan. 2018. Weaponizing the digital influence machine. *Data & Society* (2018).
- [62] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [63] Julio C. S. Reis, André Correia, Fabrício Murai, Adriano Veloso, and Fabrício Benevenuto. 2019. Supervised Learning for Fake News Detection. *IEEE Intelligent Systems* 34, 2 (2019), 76–81. <https://doi.org/10.1109/MIS.2019.2899143>
- [64] Elyse Samuels. 2020. Analysis | how misinformation on WhatsApp led to a mob killing in India. *The Washington Post* (Feb 2020). <https://www.washingtonpost.com/politics/2020/02/21/how-misinformation-whatsapp-led-deathly-mob-lynching-india/>
- [65] Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th annual meeting of the association for computational linguistics*. 1668–1678.
- [66] Shaden Shaar, Giovanni Da San Martino, Nikolay Babulkov, and Preslav Nakov. 2020. That is a Known Lie: Detecting Previously Fact-Checked Claims. *CoRR* abs/2005.06058 (2020). arXiv:2005.06058 <https://arxiv.org/abs/2005.06058>
- [67] Snopes. 2021. Transparency. *Snopes.com* (Jul 2021). <https://www.snopes.com/transparency/>
- [68] Harini Suresh and John V. Guttag. 2019. A Framework for Understanding Unintended Consequences of Machine Learning. *CoRR* abs/1901.10002 (2019). arXiv:1901.10002 <http://arxiv.org/abs/1901.10002>
- [69] Olivier Sylvain. 2018. Discriminatory Designs on User Data. *Knight First Amendment Institute at Columbia University, Forthcoming* (2018).
- [70] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for Fact Extraction and VERification. *CoRR* abs/1803.05355 (2018). arXiv:1803.05355 <http://arxiv.org/abs/1803.05355>
- [71] Thi Tran, Rohit Valecha, Paul Rad, and H Raghav Rao. 2021. An investigation of misinformation harms related to social media during two humanitarian crises. *Information systems frontiers* 23, 4 (2021), 931–939.
- [72] Kathie M. d'I. Treen, Hywel T. P. Williams, and Saffron J. O'Neill. 2020. Online misinformation about climate change. *WIREs Climate Change* 11, 5 (2020), e665. <https://doi.org/10.1002/wcc.665> arXiv:<https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcc.665>
- [73] Bertie Vidgen and Leon Derczynski. 2020. Directions in abusive language training data, a systematic review: Garbage in, garbage out. *PloS one* 15, 12 (2020), e0243300.
- [74] Nguyen Vo and Kyumin Lee. 2021. Hierarchical Multi-head Attentive Network for Evidence-aware Fake News Detection. *CoRR* abs/2102.02680 (2021). arXiv:2102.02680 <https://arxiv.org/abs/2102.02680>
- [75] Claire Wardle and Eric Singerman. 2021. Too little, too late: social media companies' failure to tackle vaccine misinformation poses a real threat. *BMJ* 372 (2021). <https://doi.org/10.1136/bmj.n26> arXiv:<https://www.bmj.com/content/372/bmj.n26.full.pdf>
- [76] John Wihbey, Matthew Kopec, and Ronald Sandler. 2021. Informational Quality Labeling on Social Media: In Defense of a Social Epistemology Strategy. (June 2021). <https://ssrn.com/abstract=3858906>
- [77] Caroline Wright, Philippa Williams, Olga Elizarova, Jennifer Dahne, Jiang Bian, Yunpeng Zhao, and Andy S L Tan. 2021. Effects of brief exposure to misinformation about e-cigarette harms on twitter: a randomised controlled experiment. *BMJ Open* 11, 9 (2021). <https://doi.org/10.1136/bmjopen-2020-045445> arXiv:<https://bmjopen.bmj.com/content/11/9/e045445.full.pdf>
- [78] Xinyi Zhou and Reza Zafarani. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. *ACM Comput. Surv.* 53, 5, Article 109 (sep 2020), 40 pages. <https://doi.org/10.1145/3395046>