XAI in Automated Fact-Checking? The Benefits Are Modest And There's No One-Explanation-Fits-All

GIONNIEVE LIM, Singapore University of Technology and Design, Singapore

SIMON T. PERRAULT, Singapore University of Technology and Design, Singapore

Fact-checking is a popular countermeasure against misinformation but the massive volume of information online has spurred active research in the automation of the task. Like expert fact-checking, it is not enough for an automated fact-checker to just be accurate, but also be able to inform and convince the user of the validity of its prediction. This becomes viable with explainable artificial intelligence (XAI). In this work, we conduct a study of XAI fact-checkers involving 180 participants to determine how users' actions towards news and their attitudes towards explanations are affected by the XAI. Our results suggest that XAI has limited effects on users' agreement with the veracity prediction of the automated fact-checker and on their intents to share news. However, XAI does nudge them towards forming uniform judgments of news veracity, thereby signaling a reliance on the explanations. We also found polarizing preferences towards XAI, raising several design considerations on these.

CCS Concepts: • Human-centered computing → Empirical studies in HCI.

Additional Key Words and Phrases: misinformation, fake news, automated fact-checking, explainable artificial intelligence (XAI), interpretable machine learning, human-centered AI, human-AI interaction

1 INTRODUCTION

Misinformation is one of the key challenges of social media. To address it, social media firms have partnered with professional fact-checkers to assess content on their platforms. These services are typically provided by independent organizations or by affiliates of news agencies and universities. Through fact-checking by experts, platform content is reviewed and rated for their accuracy and may eventually be removed if found in breach of company standards [11, 47, 55]. While effective as a measure against false information, the sheer volume of information that is created every moment is beyond the capabilities of the limited number of expert fact-checkers, making the scalability of fact-checking a fundamental concern. In seeking to address this, researchers have explored automated fact-checking. Automated fact-checking uses machine learning algorithms and artificial intelligence to assess online content and research has focused on developing various models to improve the detection accuracy of false information [38, 41, 60, 77]. There has also been research seeking to make the models more understandable to humans through the work of interpretable machine learning (IML) and explainable artificial intelligence (XAI) [79].

With ethical advances in the technology discourse urging for greater human-centered design in commercial machine learning and artificial intelligence applications, social media giants are shifting towards Responsible AI in misinformation detection. With this, it has become increasingly important to assess the effectiveness of automated fact-checkers and the practicality of explanations. In this paper, we ran an experimental study with 180 participants comparing five

Authors' addresses: Gionnieve Lim, gionnievelim@gmail.com, Singapore University of Technology and Design, Singapore; Simon T. Perrault, perrault. simon@gmail.com, Singapore University of Technology and Design, Singapore.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2023 Copyright held by the owner/author(s).

popular XAI interfaces sourced from the literature (LIME [5], SHAP [57], Attention [35], Comments [64], Evidence [54]). We sought to understand how XAI can benefit automated fact-checking in terms of the user's agreement with the fact-checker on the veracity of the news and their intent to share them. The perceptions that users had towards the XAI interfaces are also investigated. Our results suggested that while XAI interfaces have mixed effects on agreement and limited effects on the intent to share, they are relied upon to determine the veracity of news. We also found that users had polarizing preferences towards different types of XAI interfaces based on the types of visualizations and information shown. We identified five qualities that were commonly considered when participants assessed the explanations and raise several usability considerations on them.

The contribution of this work is two-fold:

- We advance efforts in social computing on the automation of the fact-checking process, particularly identifying that XAI has modest benefits by increasing the reliance on the veracity predictions provided by the automated fact-checker but has limited effects on users' overall veracity judgment and intents to share news.
- We also advance the literature on human-AI interaction by presenting a set of design considerations and a prospective user evaluation framework on explanations which, while explored in the context of fact-checking, would be relevant in other contexts given that the explanations investigated are common across other domains.

2 RELATED WORK

Our work is situated in the domains of misinformation and human-centered AI. We discuss the intersection of these domains and how XAI techniques have been proposed for fact-checking, highlighting those of interest to our study.

2.1 Misinformation on Social Media

Researchers have tracked the spread of misinformation on social media following major events like elections [6], disasters [68] and pandemics [49]. The greatest sources of misinformation may be attributed to just a few groups and people [28]. Superspreaders such as influencers, public figures, and hyper-partisan media outlets that have a large follower base post false content repeatedly and across several platforms, enabling misleading narratives to spread quickly among and across their unassuming audiences and to their cascading networks [13, 27, 39]. Social bots which are automated accounts that interact on social media in a human-like manner, also amplify misinformation through the sharing of false content, by replying to posts, and using mentions to induce greater engagement [36, 48, 69]. The rise of large language models capable of rapidly producing realistic content has also led to concerns over its potential misuse to generate false content [23].

With these contributing factors, misinformation on social media has become a key challenge of the 21st century. In response to this challenge, automated solutions have been deployed on social media platforms to detect false content and to directly remove the content or flag them for verification by professional fact-checkers [11, 47, 55]. Automated fact-checking have also been used to verify and highlight false content [26]. Yet, the complexity of these models make them difficult to understand, raising ethical concerns on their application and consequences.

2.2 Human-Centered AI

With leaping advances in AI and the application of these technologies in numerous domains, ethical concerns have been raised about the fairness, transparency and accountability of the opaque boxes [61]. Human-centered AI is a sociotechnical field seeking to amplify human agency by enhancing the human understanding of the internal workings of AI systems and empowering people to make better informed decisions. In contrast to data-centric AI that draws on statistical patterns in big data and machine learning to make decisions in ways that are incomprehensible and inaccessible to humans, human-centered AI seeks to build AI capable of explaining its decisions and is reconcilable with the moral and ethical standards of society [4].

Explainable AI is one such initiative [33]. By building a transparent model over an underlying opaque model, explainable AI aims to describe the internal mechanics of the model by explaining the features that attribute to the model's prediction. Another closely related concept often used interchangeably with XAI is IML that considers the extent to which a cause and effect can be observed directly in the underlying model. There has been active research into human-centered and explainable AI. One aspect examines how people understand, interact with, accept, and trust XAI [8, 31, 72, 78], where we place our work. XAI has been used in many domains with several benefits to Human-AI interaction and we seek to extend this investigation in the domain of misinformation that stands to benefit from explanations to assist users in understanding the automated fact-checkers' decision-making process when determining the veracity of content.

2.3 Explainable Automated Fact-checking

Social media platforms like Facebook, Twitter and TikTok have been working towards a healthier information ecosystem by incorporating deep learning technologies to their products to moderate content through flagging, labeling and removal [2, 59, 65]. In mid 2021, the companies announced their intentions towards Responsible ML and AI, in terms of both the fairness and transparency of their algorithms [3, 73]. Among their plans is to build XAI solutions that inform users on how the platform algorithms work, and the impact of the algorithms on the platform experience.

In the context of automated fact-checking that aims to predict whether a piece of content is real or fake, there is already growing research on XAI fact-checkers for media like texts, images and videos. Two popular techniques for explainability are the local interpretable model-agnostic explanations (LIME) [58] and Shapley additive explanations (SHAP) [46]. There has also been considerable interest in the attention mechanism that places weights on representations that can be visualized to provide interpretability to text-based models [7, 12, 21, 35, 56]. The attention mechanism has also been adapted in various ways to identify top social media comments that help to explain the prediction [62, 64]. In a similar vein, various modeling techniques have been used to identify top reliable evidence [50, 54].

While there are works comparing different types of explanations [66, 70], it is unclear how explanations are understood and influence users' behavior in the automated fact-checking context. Focusing on common techniques in explainable misinformation detection, we investigate how fact-check labels accompanied by the explanations that are applied to problematic posts on social media affect people's perceived veracity of news and their intents to share them.

3 METHOD

To understand the effects of different automated fact-checking explanations, we conducted a between-subjects experiment with 180 participants. Informed by prior literature [18, 42, 52], we assessed the effects by considering the participants' agreement with the fact-checker on the veracity of the news and their intent to share them.

The study sought to answer the following research questions (RQs):

- (1) What are the effects of different explanations on users' agreement with the fact-checker on the veracity of news?
- (2) What are the effects of different explanations on users' intent to share news?
- (3) What are the preferences towards explanations in automated fact-checking?

3.1 Analysis and Comparison

We investigated the most common types of XAI techniques in misinformation detection literature by systematically scanning the literature. Our search surfaced two main categories: feature-based techniques where explanations have graphical representations of weighted features (the 'LIME', 'SHAP' and 'Attention' conditions) and example-based techniques where explanations draw on textual data from other sources (the 'Comments' and 'Evidence' conditions) (Figure 1). To compare our results, we also included a *Control* condition where no XAI explanation is provided.



Fig. 1. The XAI conditions for the fact-checker report based on the TRUE news shown in Figure 2, each with a "How to understand this?" guide. The first three conditions (LIME, SHAP, Attention) display similar data with a graphical visualization, while Evidence and Comments use a textual visualization.

We assessed the main effects of the XAI techniques by analysing all five XAI conditions and the control. For a finer discussion, we focused on the comparison between techniques of the same graphical or textual categories as comparing across categories would be less meaningful. For instance, we can be confident that the reason was due to the difference in the visualization for a difference between LIME and SHAP (both graphical), whereas if the difference was between LIME (graphical) and Evidence (textual), it would be difficult to tell whether it was due to the visualization or the information presented.

3.2 Procedure

We collected the data using an online webapp created by the researchers. Participants were recruited on Amazon Mechanical Turk (MTurk) through Human Intelligence Tasks (HITs). Participation was voluntary and it was possible to withdraw at any time with no data being collected. The study was approved by the University's Institutional Review Board (IRB) and participants were reimbursed for their time at a rate of 7.25 USD/hour as per our IRB guidelines.

When participants accepted the HIT, they were given a link to the webapp to complete a screener survey. They would be redirected to the study upon passing. There were 12 links for the random redirection, six for each XAI condition and the control, and another six with counterbalancing of the Likert scales. Participants who did not pass were advised to return the HIT instead. After filling their demographic details, participants proceeded to the experiment. A walkthrough containing descriptions of the news, automated fact-checker, and the task instructions were provided. Participants then proceeded to the task of rating a randomized set of 12 news with the assigned XAI interface. For each news, they provided their perceived veracity of the news (4-points, False to True) and their intent to share it (5-points, Definitely Not to Definitely Yes). Thereafter, they completed a post-experiment survey containing an open-ended question on their perceptions towards the assigned XAI interface. Upon completion, participants were given instructions to generate a unique completion code to submit on MTurk as a proof of completion.

3.3 Materials

3.3.1 Stimuli/News. For the study, 12 news on political topics were used with an equal ratio of true to false news. The news were sourced from Snopes.com, a fact-checking service that uses an extensive rating system to cover the various dimensions of misinformation, within three months of the study. We only used news that Snopes rated to be either 'True' or 'False' [63]. We then searched for the original news headlines from a variety of news websites and made grammar and case modifications to the headlines for a consistent reporting style. To prevent the source of news from influencing credibility perceptions of the content [74], we chose a plausible yet neutral-sounding source domain for all the news: 'newsarchives.com', that at the time of the study was not an active site. An interface similar to the Facebook news feed post that includes a label for the veracity prediction of the automated fact-checker was used for the news (Figure 2). This also served as the control.



Fig. 2. News that received a TRUE or FALSE prediction by the automated fact-checker. These visualizations were used as Control for our experiment.

3.3.2 XAI Interfaces. Participants were shown the XAI interface based on their assigned condition. The five XAI conditions were: LIME, SHAP, Attention, Comments, and Evidence. Each XAI interface contained a news (Figure 2) and a fact-checker report with a "How to understand this?" guide describing various aspects of the explanations (Figure 1). Larger individual images of the XAI interfaces can be found in Appendix A. With the variety of conditions, we carefully considered the design of the fact-checker reports to ensure homogeneity in them and created mockups of the

fact-checker reports instead of building actual XAI models to do so. Due to the marked differences in the visualization and mechanism of the conditions, we designed for homogeneity across LIME ([5] Figure 14), SHAP ([57] Figure 7) and Attention ([35] Table 5), and across Comments ([64] Figure 1) and Evidence ([54] Figure 2). We aimed to follow the original design of each condition (as referenced) as closely as possible while doing so.

As explained in the 3.1, we will discuss differences across XAI interfaces from the same category (*Graphical* vs. *Textual*) as comparison across categories are unfair.

Selection of conditions. To identify relevant papers, keyword searches were conducted in the ACM Digital Library and IEEE Xplore, the two largest computing literature databases. A three-part query using combinations of {'false information', 'misinformation', 'disinformation', 'fake news'} AND {'explain*', 'interpret*'} AND {'artificial intelligence', 'AI', 'machine learning', 'ML'} returned 308 ACM and 1061 IEEE unique results on 10 July, 2021. Screening of the results were conducted with the criteria that the articles were (1) full papers, (2) about an XAI or IML technique, (3) for misinformation detection, and (4) using texts. We found 31 papers that fit the criteria. A thorough reading of them identified the attention mechanism [7, 12, 21, 56], evidence from other news sources [50, 53, 54, 71], and SHAP [29, 57, 67] as more popularly studied forms of explanations. Other techniques like knowledge graphs, decision trees and syntactic patterns had only one to two relevant papers each. Additionally, LIME [5, 40] was selected due to its generic popularity in XAI while explanations using social media comments [22, 64] was included due to its similarity to evidence from other news sources. This led to the selection of the five XAI conditions.

Graphical XAI Interfaces. For these interfaces, the automated fact-checker makes its prediction by weighing each unit of word in a line of text and providing a graphical representation of the contribution of each word as an explanation. For homogeneity, we designed all three XAI consider only the words in the news. In doing so, we can make it such that the scores and the effects of the words were the same. For example, if the 'fake' prediction probability in LIME is 0.26, the percentage of red arrows in the arrow bar in SHAP would be 26%, and Attention would have a prediction score of 0.26. The effect of the words is treated in a similar fashion where words with greater effect have deeper shades or longer bars. As such, these conditions are also content-based as they explain using the news.

Textual XAI Interfaces. For the second category, the automated fact-checker makes its prediction by weighing information from secondary sources, i.e., social media feeds for Comments and other news articles for Evidence, and using the top matches as explanations. Since we used real comments from Twitter and Reddit and real articles from a variety of news websites to more truthfully emulate the XAI, it was largely impossible to get comments and evidence that were in parallel to each other. Therefore, for homogeneity in design we focused on the visual aspect of the explanations by using the same layout ([64] figure 1): a description at the top followed by three boxes that each have a score at the bottom right corner, using similar color schemes to indicate the stance of the boxes, and always having a box that either supports or is neutral towards the news as the top-most one. As such, these conditions are also context-based, as they explain via alternative information surrounding the news headline.

3.4 Variables

Independent Variables. The first independent variable is the XAI condition (XAI CONDITION). This independent variable is between-subjects: participants are exposed to only the 'LIME', 'SHAP', 'Attention', 'Comments' or 'Evidence' condition (see Figure 1). We also included 'Control' that only includes an indication of whether the headline is true or false (see Figure 2). All five XAI conditions are described in section 3.3.2.

The second independent variable is the veracity of the news (NEWS VERACITY) and is within-subject with each has being either 'True' or 'False'. We assess this variable as an interaction effect with XAI CONDITION.

Dependent Variables. The first dependent variable is the agreement with the fact-checker (*Agreement*). Participants were asked "How accurate do you think the news above is?" on a 4-point Likert scale {1: False, 2: Somewhat False, 3: Somewhat True, 4: True}. This did not have a neutral statement as we wanted the participants to take a decision on the veracity of the news, following prior literature assessing misinformation interventions [18, 42, 52]. We then converted the perceived veracity into an agreement score by comparing it to the actual veracity of the news. A participant answering "True" on a real news (resp. "False" on a fake news) would reach an agreement score of 4, "Somewhat True" (resp. "Somewhat False") a score of 3, and so forth. We also compared the variances of the agreement score as a proxy to measure the reliance of users on the XAI when forming a veracity judgment.

The second dependent variable is the intent to share the news through "I would consider sharing this news on social media." (*Intent to Share*). This variable has been examined in earlier literature due to its direct influence on the spread of misinformation [18, 42, 52] and is asked on a 5-point Likert scale {1: Definitely Not, 2: Somewhat Not, 3: Neutral, 4: Somewhat Yes, 5: Definitely Yes}.

3.5 Design and Iterations of the Study

Through several pilot studies, we refined the study by rephrasing the questions for better clarity, adding a progress indicator, and reducing the number of news shown. Realizing that participation fatigue was a major concern in our design due to the repetitiveness of the task and the cognitive burden in processing the fact-checker reports, we reduced from an initial set of 24 to 12 news. This reduction allowed for a more manageable completion time of around 15 minutes. While this led to less observations, we considered the trade-offs such as abandonment of the experiment or lower quality responses in later trials due to boredom or distraction [30] and agreed on the reduction.

Overall, we gathered a total of 30 (participants) × 6 (XAI conditions) × 12 (news) = 2160 trials.

3.6 Participants and Recruitment

In reviewing the collected data, the exclusion criteria included (1) failing the screener survey, (2) failing the attention check, (3) completing the study within an exceptionally short or long time, (4) responding randomly during the study, and (5) searching online for more information during the study, Criteria 4 and 5 were assessed through respective Yes/No questions set at the end of the study. Of the 222 responses received, 42 were excluded for the above reasons.

The study had 180 participants. Their age ranged from 23 to 73 years old (M = 38.4, SD = 10.5) and a majority identified as male (57.8%). On political ideology, a majority identified as democrats (62.8%). On education, a majority had Bachelor's degrees (55.0%). Through 5-point Likert scale questions (1: Definitely Not, 5: Definitely Yes) on news reading habits, participants reported reading news on social media (M=3.96, SD=1.10) and traditional news media (M=3.81, SD=1.29), and following fact-checked content (M=3.41, SD=1.15).

We adopted a rolling recruitment process until 30 accepted participants per condition were reached. On MTurk, workers had to be located in the US, with a HIT approval rate of at least 98%, and more than 5000 approved HITs. Custom qualifications were used to prevent repeated attempts. Through the screener survey, participants were assessed to (1) be 18 years old and above, (2) live in the US for the last 5 years, (3) be fluent in English, (4) actively use and (5) share content on social media, and (6) agree to participate.

3.7 Data Analysis

Statistical Analysis. Since the data was not normally distributed, a non-parametric approach was used for the data analysis. For XAI CONDITION, Kruskal-Wallis tests were used for main effects followed by pairwise Mann-Whitney U tests for pairwise comparisons. Aligned rank tests [75] was used for interaction effects between XAI CONDITION and NEWS VERACITY followed by contrast tests [24] for pairwise comparisons. Finally, we used Fligner-Killeen's test to compare the variance of the *Agreement* scores across the XAI conditions and used the same for the pairwise comparisons. Benjamini-Hochberg corrections were used for all pairwise comparisons to minimize false discovery rates [32].

Qualitative Coding for the Post-Experiment Survey. We conducted focused coding [25] on participants' responses to their perceptions of the assigned XAI interface. Through initial rounds of open coding, a codebook (Table 2) was developed which was subsequently applied to the responses to capture common threads in them. Training of the coders occurred over three iterations, where disagreements were resolved through discussion and the definitions and protocols of the codebooks were refined. Two coders were trained on a random sample of 20% of the responses, reaching an average interrater reliability of $\kappa = .89$ (SD = .11, Min = .75, Max = 1), using Cohen's Kappa [19].

4 RESULTS: EXPERIMENT

4.1 Addressing RQ1: Effects of Explanations on User's Agreement with the Fact-Checker

There was a main effect of XAI CONDITION on Agreement ($\chi^2(5) = 14.58$, p = .012, $\eta^2 = .027$). Participants in the SHAP condition (M = 3.49, SD = 0.44) had significantly higher agreement scores than the Control (M = 3.24, SD = 0.57) as well as the Attention (M = 3.21, SD = 0.53, all p < .05) conditions. We also found a significant difference between SHAP and Evidence (M = 3.31, SD = 0.41) but as discussed in section 3.1, comparisons across XAI categories can be problematic. Results are shown in Figure 3a.

The variances of *Agreement* across XAI CONDITION were not homogeneous ($\chi^2(3) = 12.13, p = .033$) as shown in Figure 3b. Both LIME (*SD* = 0.39, *IQR* = .33) and SHAP (*SD* = 0.44, *IQR* = .50) had significantly lower variance (both p < .05) to the Control (*SD* = 0.57, *IQR* = .67). No difference between the *Textual* XAI interfaces and control were found.



Fig. 3. The left chart shows the effects of XAI Condition on Agreement. Agreement scores are between 1 (strong disagreement) and 4 (strong agreement) with the averages shown. The right chart shows the variance in the agreement scores for each condition. Error bars show .95 confidence intervals. *: p < .05, **: p < .001, ***: p < .001.

4.1.1 Interaction Effects on Agreement. We found a significant XAI CONDITION × NEWS VERACITY interaction on Agreement ($F_{5,174} = 4.13$, p < .01, $\eta^2 = 0.12$). Participants in the Attention condition had significantly lower agreement on true news than false news (p = 0.01). There were also an interesting observation with Evidence where the agreement was higher for true news than false news compared to all the other conditions where the agreement for false news were higher than true news.



Fig. 4. Interaction of XAI Condition and News Veracity on Agreement. Scores are between 1 (strong disagreement) and 4 (strong agreement) with the averages shown. Error bars show .95 confidence intervals.

4.2 Addressing RQ2: Effects of Explanations on User's Intent to Share News

No main effect of XAI CONDITION on the *Intent to Share* was found ($\chi^2(5) = 5.50$, p = .36, $\eta^2 = .0014$). While the LIME, SHAP and Comments conditions seemed to lower the intent to share, this was not observed for the Attention and Evidence conditions. The descriptive statistics for each XAI condition is shown in Table 1. We also did not find any interaction effects with news veracity (p = .25).

Table 1. Descriptive statistics (M(SD)) of the XAI conditions on the intent to share. Scores are between 1 (strongly disagree) and 5 (strongly agree).

	Control	LIME	SHAP	Attention	Comments	Evidence
Share	2.28 (1.00)	2.08 (1.02)	2.04 (1.04)	2.29 (1.00)	2.13 (0.87)	2.29 (0.92)

5 RESULTS: POST-EXPERIMENT SURVEY

We describe the results of the post-experiment survey, addressing RQ3: What are the preferences towards explanations in automated fact-checking?

5.1 Perceptions towards the Assigned XAI Interfaces

For the open-ended question on the opinions on the XAI condition assigned, we excluded the Control condition, thereby analyzing 150 responses by participants in the five XAI conditions. Of these, 17 responses that were empty, not meaningful or irrelevant were discarded. In total, 133 responses with an average of 122.7 characters (SD = 84.2, Min = 6, Max = 396)

Code	Subcode	Frequency (<i>n</i>)	Frequency (%)	IRR (κ)				
	Positive	69	51.9					
Continuent	Negative	36	27.1	.75				
Sentiment	Conditional	7	5.3					
	Mixed	21	15.8					
Qualities of the XAI Fact-Checker								
	Positive	46	34.6	.93				
Usefulness	Negative	9	6.8					
	Conditional	4	3.0					
	Positive	23	17.3					
Ease of Understanding	Negative	14	10.5	.88				
	Conditional	0	0.0					
	Positive	26	19.5					
Accuracy	Negative	9	6.8	1.0				
	Conditional	2	1.5					
	Positive	7	5.3					
Trustworthiness	Negative	23	17,3	.75				
	Conditional	4	3.0					
	Positive	19	14.3					
Design	Negative	8	6.0	1.0				
_	Conditional	0	0.0					

Table 2. Codebook on the opinions on the assigned XAI condition (N = 133). There are 6 code categories and 19 subcodes with the frequencies and interrater reliability values reported.

were analyzed. A summary of the codebook and results are shown in Table 2. We coded the general sentiment and five qualities of the XAI fact-checker: usefulness, ease of understanding, accuracy, trustworthiness and design. The categories are further split into positive, negative, conditional and mixed. Conditional refers to aspects that will be fulfilled only when some condition is satisfied and is thus neither positive nor negative. Mixed refers to aspects containing combinations of positive, negative and conditional elements.

The qualities are described as follows:

- Usefulness Level of aid provided by the fact-checker to help users ascertain the veracity of news
- Ease of Understanding Effort required by users to interpret and comprehend the explanations
- Accuracy Performance of the fact-checker in consistently making a correct decision on the veracity of news
- Trustworthiness Reliability of the features used to build the algorithm and the ownership of the fact-checker
- Design Visual appeal and organization of the explanations

5.2 General Positive Opinions

Participants displayed a largely positive opinion for the different XAI conditions as reported in Table 2. Forty-six participants found the automated fact-checker useful or helpful with some participants noting that beyond fact-checking, the explanations made them consider the news more critically and to be more cautious. P144 mentioned that the fact-checker *"helps me look at news stories with an open-mind, and with less bias. I'm less likely to assume something is true,*

only because I want it to be true". P164 further commented on autonomy saying that "I like that it gives people a way to determine if something is true or false, letting the person make that decision based on the facts given." For the 23 participants who found the fact-checker easy to understand, they described it as "intuitive" (P105), "straight forward" (P175), and "clear and concise and very easy to read" (P70). Twenty-six participants found the fact-checker accurate and seven found it trustworthy with P139 saying that it "seems very reliable and I agreed with most of what the checker had determined to be real versus fake." These aspects could have been evaluated through the participants' personal understanding of the news as mentioned by P147 who said that "I recognized a few of the news stories and felt [that] the fact-checker was correct in it's assessment of them." For the 19 participants who positively commented on the design of the fact-checker, they described it as "attractive" (P44), "simple and clean" (P38), with a "well structured format" (P98) and "colors [that] were cool to look at" (P149). In addition, several participants suggested that they would want to have the XAI fact-checker in their everyday life. As put by P176, "I think that tools like this one will become increasingly necessary in the age of social media and rapid dissemination of misinformation."

5.3 General Negative Opinions

Most of the negative opinions centered on the trustworthiness of the fact-checker with 23 participants finding it untrustworthy. First, there were doubts on AI technology itself. P154 said "I am always a bit 'leery[]' with any information that is generated by an algorithm, because the algorithms are made by humans, and humans are fallible". Second, the suitability of automating the fact-checking process was questioned with P49 saying that "I do not think [AI] fact checkers should ever[] be use[d] as they simply cannot understand context or be programmed in a way that can reliably fact check sources." Lastly, a common concern with several participants was the manipulation of fact-checkers (both automated and not) to serve political purposes. Fact-checkers were thought to be "heavily biased propaganda machines meant to push a narrative" (P48). Strong defiance was expressed in these cases, with P108 asking "Who is fact checking the fact checkers[?] I can make up my own mind and don[']t need communist [sic] telling me what to think."

As for the other qualities, nine participants felt that the fact-checker was not useful or needed. For P57, the factchecker "created more doubts for me when it was in the middle of true or false!! [sic]" while P1 said that "I investigate any news that is interesting to me on my own. I really do not need a fact checker." Fourteen participants had difficulty understanding the fact-checker, finding the explanations confusing. P9 stated that "I would likely do additional research on articles before I share them at least 75% of the time" as a result. The nine participants who found the fact-checker inaccurate noted that there were errors made by the fact-checker such as there being "some things that were off with some of the stories that were definitely true" (P153). As for the eight participants who commented negatively on the design, they felt that the explanation was "cumbersome" (P92), "very cluttered [and] needs to be simplified" (P96).

5.4 Opinions by the XAI Condition

Figure 5 shows the breakdown of the qualities mentioned for each assigned XAI condition. We further provide some quotes for each condition.

For LIME, there were contradicting opinions with some participants like P96 commenting that "it was very cluttered needs to be simplified" yet others like P118 saying that it had a "Simple and easy to use design.". The higher number of visuals could also have led P30 to comment that "I think that it's accurate, but I find the way it explains how it comes to its conclusions to be a bit confusing."

There were similar observations for SHAP with P159 saying that "I did not fully understand how the fact checker mad[e] its determinations" yet P29 commenting that "I thought it was a good tool and the graph was helpful in understanding its

conclusions." These was likely due to a learning curve in understanding the explanations. As put by P3, "I was a bit confused by how the fact checker worked, but was able to decipher it as the task went along".

For Attention, some positive comments included finding it "simple" (P67) with a design that was "pretty good and easy to follow" (P133) whereas some noted that it "did not highlight some important keywords in the news" (P152) and that "the prediction score needs to be worked on as it can be confusing" (P167).

For Comment, P67 said that "It's graphic design is so-so but the content is trustworthy and compelling." although P119 instead commented that "I am under the impression that the fact-checking comes from the comments? Which seem completely biased and unreliable.. makes no sense!"

For Evidence, P79 "I like the layout and ease of use of it; it prompted me to think about the content I was consuming and made me question the articles." and that "The snippets it gave to read were helpful." (P90). However, there were also concerns about the source with P14 commenting that "This is a difficult question in a divided nation where propaganda, lies, and innuendo govern any "news" produced by the conservative right."

In general, there were polarising perceptions towards the XAI conditions in both their graphical and textual design and to the content and contextual information presented for the fact-checking.



Fig. 5. Breakdown of the qualities mentioned for each XAI condition.

6 DISCUSSION

6.1 Differences in Agreement Between the XAI Conditions

The XAI conditions showed mixed results on the level of agreement with the fact-checker.

Comparing to the Control. Of all the XAI conditions, only SHAP led to significantly higher levels of agreement. This suggests that the presence of explanations itself may not strongly convince users of the automated fact-checkers's prediction.

Between Graphical Explanations: LIME, SHAP and Attention. The Attention condition had lower levels of agreement to LIME and SHAP even though all three conditions were designed to give the same output. The additional dimension of the polarity of the features in LIME and SHAP could have been more visually impactful and thereby more convincing compared to Attention which only had a single shade. As such, apart from the algorithm used to perform the fact-checking, the way the results are explained to users is also important as design changes may lead to weaker persuasiveness of the XAI, and consequently, perhaps even lower levels of trust in them [51].

Between Textual Explanations: Comments and Evidence. From the interaction effect with News Veracity, Evidence had higher agreement scores for true news and lower for false news, whereas Comments had higher scores for both, compared to the Control. The observation with Evidence on false news is rather surprising. Were participants unconvinced by the excerpts from news articles? Or worse, were they instead convinced of the otherwise by them? Perhaps the weakening trust in media outlets is at play here [37]. If so, this is a concerning observation that calls for a broader investigation. In contrast, while social media is not considered authoritative in the traditional sense, participants do seem to take to the content. This may be because social media comments tend to be more opinionated and written with expressive and emotional language that thereby come across as more persuasive.

6.2 XAI Reduces Variance in Agreement

Our analysis on the variance of the agreement scores showed that the variances of the five XAI conditions were lower than the control despite their comparable mean scores, with the LIME and SHAP conditions being significantly so, possibly because they share similar features [70]. A larger spread means a greater variation in the levels of agreement within a group of participants, and as such, the lower observed variance suggests that explanations, particularly LIME and SHAP, can help users determine whether a news is true or not more confidently. This is in line with the work of Lai and Tan [43] that found that 'additional details including irrelevant ones can improve the trust that humans place on machine predictions' where trust was defined as 'the percentage of instances for which humans follow the machine prediction'. This demonstrates a clear benefit of explanations in that they are relied upon for determining the veracity of news and that they help in reinforcing veracity judgments, thereby leading to greater uniformity in the veracity perceptions.

6.3 XAI Has Limited Impact on the Intent to Share

As we did not observe any effect of the XAI condition on participants' intent to share, we conclude that the different conditions we chose for the study do not seem to affect participants' behavior in terms of sharing habits. This result suggests that automated fact-checkers may help users assess the veracity of news, but that incorporating XAI on a social network platform would not directly impact the number of shares of a given piece of news. Studies have shown that social factors are stronger motivators for sharing [16]. Thus, our observation is aligned with previous research.

6.4 No Best Type of Explanation

There were polarizing preferences towards the graphical/content-based and textual/context-based XAI interfaces among the participants from the post-experiment survey. This suggests that there is no perfect form of explanation for

automated fact-checking, echoing the findings of work in other domains [10, 34, 76]. One's background, knowledge and experience shape one's being and this leads to the diversity of subjective preferences in people. This makes it challenging to settle on a single explanation as has been the case for user interface design in general [15] which has led to the exploration of alternatives such as adaptive interfaces [14]. More investigations on examining and identifying explanations that are suitable or 'good enough' for the contexts and goals of their use will be necessary.

7 IMPLICATIONS FOR DESIGN

7.1 Usability Considerations

From the study, we surface some observations on the usability of XAI interfaces. First, content and visualization matters as incomprehensible or unconvincing explanations may lead to effects that are lower than or even opposite to what is expected. Second, with the widely varying preferences, providing a variety of explanations to enable users to have a comprehensive overview [17] or tailoring an explanation to suit the needs of the user [45] would be more adequate than a single standardized explanation. Third, for the context-based mechanisms, participants displayed greater trust in evidence and greater distrust in comments, thereby mirroring the same impression onto the fact-checker. With explanations revealing the features of the AI, more tactfulness should be taken in the selection of the sources and factors used in the algorithm. Fourth, with information processing being cognitively burdensome, the amount of explanations provided by the fact-checker should be kept down. This was exemplified by how we had to reduce the number of news presented to participants as a result of cognitive taxation. Efforts should be taken to design concise explanations that minimize cognitive load [1] which is not a trivial task as explanations also have to be informative. More research in this area would be beneficial towards the public adoption of XAI in any field. Lastly, and in line with the previous point, the integration of explanations in social media apps should not clutter the interface, particularly on mobile phones that have limited screen space. A plausible design suggested by P177 is *"I'd rather something small but eye catching that says true or false, or mixed, and ... a link that explains rather than such a large explanation under the article."*

7.2 Prospective User Evaluation Framework

From the open-ended questions, we identified five qualities that were consistently brought up in the feedback. These are the usefulness, ease of understanding, accuracy, trustworthiness and design of the fact-checker. While our study was not centered on the development of a framework, we realized that these qualities revealed aspects of the XAI that mattered most strongly to users. As such, we posit that they may factor as a framework in the user evaluation of XAI. Such a framework would be beneficial in assessing explanations in general, such as in other highly consequential domains of healthcare and law, where XAI is used to inform users including professionals and the general public, of the decisions made by the algorithms. By having user-generated principals to assess explanations, this aligns with and supports the human-centered design of XAI. Langer et al. have listed 29 stakeholders' desiderata on explainability [44], some of which is in common with those we describe, namely Usefulness, Accuracy and Trustworthiness, and there are others like Confidence and Satisfaction where parallels may be drawn with the remaining qualities. While the qualities we identify are not nearly as comprehensive as theirs, we believe that these qualities are intuitively more important to users and can serve as a concise user evaluation framework.

In the interest of continuity, we offer a brief description of the qualities:

- Usefulness Level of aid provided by the XAI to help users meet a practical goal
- Ease of Understanding Effort required by users to interpret and comprehend the explanations

- · Accuracy Performance of the XAI in consistently making a correct decision
- Trustworthiness Reliability of the features used to build the algorithm and the ownership of the XAI
- Design Visual appeal and organization of the explanations

This prospective framework is, however, a coarse one that demands further investigation.

7.3 Reflecting on Explainable Automated Fact-Checking

During the qualitative coding process, we encountered two statements that made us contemplate on the practicality of using explanations in automated fact-checking. The first was by P52 who said that "text analysis can be gamed" and the second by P12 saying that "I'd be worried of its veracity [in the] long-term, especially if people find out how it works and can then purposefully word their posts to game the system." Both participants were speaking of content-based explanations and they were concerned that over time, users could easily figure out how the fact-checker weighted words, thereby crafting their posts in favor of a positive veracity and devaluing the automated fact-checker. We felt that this was a critical observation and sought to have an expanded discussion on it.

With explanations exposing the decision-making process of the automated fact-checker, it also poses a risk for malicious actors to unravel and attempt to outwit the system. Such concerns have been echoed before with Twitter saying that *"to ensure people cannot circumvent these safeguards, we're unable to share the details of these internal signals in our public API"* for their bot detection algorithm [20]. This also raises the question of whether explanations are ready to be released to the public even if they are well developed, especially in the fields of surveillance, security and safety. And if so, what types of data and XAI techniques are suitable for public-facing systems that hold such risks? For explainable automated fact-checking, perhaps example-based or generated natural language explanations [9] are best whereas feature-based explanations that reveal the inner workings of the algorithm should be kept for internal use such as by the system architects and the fact-checking experts engaged to review detected misinformation.

8 LIMITATIONS

The study was run on MTurk with only US-based participants. As such, results may not generalize broadly to different countries and contexts. Design-wise, the study was done with the XAI condition as a between-subjects factor in consideration of the study duration. Using a within-subjects design for this variable could have led to better comparisons but would likely have induced greater participation fatigue. Also, as we examined only five types of XAI conditions widely used in the literature, further studies may want to investigate new designs or existing ones that were not included.

The study was done while explainable automated fact-checkers are not in the mainstream which might raise doubts on its applicability. Yet, there are already indicators of eager advances to use automated fact-checking to aid in the work of human fact-checkers and to potentially expose to users with the turn towards Responsible AI by social media firms. The perceptions of AI and fact-checkers may continue to evolve in the future and this work paves a path towards that understanding by investigating explanations that are at the fore of current explainable misinformation detection research. Furthermore, beyond the context of fact-checking, several of the considerations surfaced in this study are on the explanations themselves, rather than the context they are applied in, which are generic and applicable in other domains.

9 CONCLUSION

To understand the behavioral effects of explanations and how they are perceived by users, we conducted a study on five XAI interfaces in the field of misinformation where efforts have been made to improve the explainability of automated fact-checkers. We found various effects on users' veracity perception of news and their intents to share and check them by the XAI condition and the veracity of the news. In particular, we observed that XAI has mixed or limited effects on the agreement with the fact-check and the intent to share news, two aspects that are important performance indicators on the effectiveness of fact-checking. Despite so, there is a reliance on explanations in the veracity judgments of news as shown by the decrease in agreement variances when compared to the control. As such, we conclude with the first half of our title: while not in the key areas, XAI does offer modest benefits in automated fact-checking. Our participants also displayed polarizing preferences between the graphical/content-based and textual/context-based explanations, with some raising concerns on the effectiveness, reliability and usability of the XAI fact-checkers. Thereby, we conclude with the second half of our title: that there is no single type of explanation that suits all. In light of this, we consolidated several considerations and suggestions on making explanations more practicable in both the context of fact-checking and in general. We also put forward a prospective user evaluation framework for explanations that was formulated from participants' feedback on XAI interfaces. With explainability coming back to the fore in recent years, we aim for this work to contribute to more human-centered practices in the design of XAI. As future work, we plan to conduct more studies to understand the effects we have observed and to characterize the user types for the varying explanation preferences.

REFERENCES

- Ashraf Abdul, Christian von der Weth, Mohan Kankanhalli, and Brian Y. Lim. 2020. COGAM: Measuring and Moderating Cognitive Load in Machine Learning Model Explanations. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/3313831.3376615
- [2] Facebook AI. 2020. Here's how we're using AI to help detect misinformation. Retrieved June 1, 2023 from https://ai.facebook.com/blog/heres-howwere-using-ai-to-help-detect-misinformation/
- [3] Facebook AI. 2021. Facebook's five pillars of Responsible AI. Retrieved June 1, 2023 from https://ai.facebook.com/blog/facebooks-five-pillars-of-responsible-ai/
- [4] AI4EU. 2020. Human-Centred AI: What is going on in Europe. Retrieved June 1, 2023 from https://www.ai4eu.eu/news/human-centred-ai-whatgoing-europe
- [5] Hassan Ali, Muhammad Suleman Khan, Amer Alghadhban, Meshari Alazmi, Ahmad Alzamil, Khaled Al-Utaibi, and Junaid Qadir. 2021. All Your Fake Detector are Belong to Us: Evaluating Adversarial Robustness of Fake-News Detectors Under Black-Box Settings. *IEEE Access* 9 (2021), 81678–81692. https://doi.org/10.1109/ACCESS.2021.3085875
- [6] Hunt Allcott, Matthew Gentzkow, and Chuan Yu. 2019. Trends in the diffusion of misinformation on social media. Research & Politics 6, 2 (2019), 2053168019848554. https://doi.org/10.1177/2053168019848554 arXiv:https://doi.org/10.1177/2053168019848554
- [7] Nujud Aloshban. 2020. ACT : Automatic Fake News Classification Through Self-Attention. In 12th ACM Conference on Web Science (Southampton, United Kingdom) (WebSci '20). Association for Computing Machinery, New York, NY, USA, 115–124. https://doi.org/10.1145/3394231.3397901
- [8] Ariful Islam Anik and Andrea Bunt. 2021. Data-Centric Explanations: Explaining Training Data of Machine Learning Systems to Promote Transparency. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21). Association for Computing Machinery, New York, NY, USA, Article 75, 13 pages. https://doi.org/10.1145/3411764.3445736
- [9] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. Generating Fact Checking Explanations. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, Online, 7352–7364. https: //doi.org/10.18653/v1/2020.acl-main.656
- [10] Daniel Ben David, Yehezkel S. Resheff, and Talia Tron. 2021. Explainable AI and Adoption of Financial Algorithmic Advisors: An Experimental Study. In Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society (Virtual Event, USA) (AIES '21). Association for Computing Machinery, New York, NY, USA, 390–400. https://doi.org/10.1145/3461702.3462565
- [11] Arjun Narayan Bettadapur. 2020. TikTok partners with fact-checking experts to combat misinformation. Retrieved June 1, 2023 from https: //newsroom.tiktok.com/en-au/tiktok-partners-with-fact-checking-experts-to-combat-misinformation
- [12] Luís Borges, Bruno Martins, and Pável Calado. 2019. Combining Similarity Features and Deep Representation Learning for Stance Detection in the Context of Checking Fake News. J. Data and Information Quality 11, 3, Article 14 (June 2019), 26 pages. https://doi.org/10.1145/3287763

Behavioral Effects and User Preferences on XAI Fact-Checkers

- [13] Alexandre Bovet and Hernán A. Makse. 2019. Influence of fake news in Twitter during the 2016 US presidential election. Nature Communications 10, 7 (Jan. 2019). https://doi.org/10.1038/s41467-018-07761-2
- [14] Dermot Browne, Peter Totterdell, and Mike Norman. 1990. Adaptive User Interfaces. Academic Press, London.
- [15] John M. Carroll. 2003. Interfacing Thought: Cognitive Aspects of Human-Computer Interaction. The MIT Press.
- [16] Xinran Chen and Sei-Ching Joanna Sin. 2013. 'Misinformation? What of It?': Motivations and Individual Differences in Misinformation Sharing on Social Media. In Proceedings of the 76th ASIS&T Annual Meeting: Beyond the Cloud: Rethinking Information Boundaries (Montreal, Quebec, Canada) (ASIST '13). American Society for Information Science, USA, Article 104, 4 pages.
- [17] Michael Chromik and Andreas Butz. 2021. Human-XAI Interaction: A Review and Design Principles for Explanation User Interfaces. In Human-Computer Interaction – INTERACT 2021, Carmelo Ardito, Rosa Lanzilotti, Alessio Malizia, Helen Petrie, Antonio Piccinno, Giuseppe Desolda, and Kori Inkpen (Eds.). Springer International Publishing, Cham, 619–640.
- [18] Katherine Clayton, Spencer Blair, Jonathan A. Busam, Samuel Forstner, John Glance, Guy Green, Anna Kawata, Akhila Kovvuri, Jonathan Martin, Evan Morgan, Morgan Sandhu, Rachel Sang, Rachel Scholz-Bright, Austin T. Welch, Andrew G. Wolff, Amanda Zhou, and Brendan Nyhan. 2020. Real Solutions for Fake News? Measuring the Effectiveness of General Warnings and Fact-Check Tags in Reducing Belief in False Stories on Social Media. *Political Behavior* 42 (2020), 1073–1095. https://doi.org/10.1007/s11109-019-09533-0
- [19] Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. Educational and Psychological Measurement 20, 1 (1960), 37–46. https://doi.org/10.1177/001316446002000104 arXiv:https://doi.org/10.1177/001316446002000104
- [20] Colin Crowell. 2017. Our approach to bots and misinformation. Retrieved June 1, 2023 from https://blog.twitter.com/en_us/topics/company/2017/Our-Approach-Bots-Misinformation
- [21] André Ferreira Cruz, Gil Rocha, and Henrique Lopes Cardoso. 2020. On Document Representations for Detection of Biased News Articles. In Proceedings of the 35th Annual ACM Symposium on Applied Computing (Brno, Czech Republic) (SAC '20). Association for Computing Machinery, New York, NY, USA, 892–899. https://doi.org/10.1145/3341105.3374025
- [22] Limeng Cui, Kai Shu, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. DEFEND: A System for Explainable Fake News Detection. In Proceedings of the 28th ACM International Conference on Information and Knowledge Management (Beijing, China) (CIKM '19). Association for Computing Machinery, New York, NY, USA, 2961–2964. https://doi.org/10.1145/3357384.3357862
- [23] Luigi De Angelis, Francesco Baglivo, Guglielmo Arzilli, Gaetano Pierpaolo Privitera, Paolo Ferragina, Alberto Eugenio Tozzi, and Caterina Rizzo. 2023. ChatGPT and the rise of large language models: the new AI-driven infodemic threat in public health. Frontiers in Public Health 11 (2023). https://doi.org/10.3389/fpubh.2023.1166120
- [24] Lisa A. Elkin, Matthew Kay, James J. Higgins, and Jacob O. Wobbrock. 2021. An Aligned Rank Transform Procedure for Multifactor Contrast Tests. In The 34th Annual ACM Symposium on User Interface Software and Technology (Virtual Event, USA) (UIST '21). Association for Computing Machinery, New York, NY, USA, 754–768. https://doi.org/10.1145/3472749.3474784
- [25] Kristin G. Esterberg. 2002. Qualitative methods in social research. McGraw-Hill.
- [26] Factinsect. [n.d.]. Factinsect, the Automated AI Fact Checker. Retrieved June 1, 2023 from https://factinsect.com/
- [27] Center for an Informed Public, Digital Forensic Research Lab, Graphika, and Stanford Internet Observatory. 2021. The Long Fuse: Misinformation and the 2020 Election. Retrieved June 1, 2023 from https://purl.stanford.edu/tr171zs0069
- [28] Knight Foundation. 2018. Disinformation, 'Fake News' and Influence Campaigns on Twitter. Retrieved June 1, 2023 from https://knightfoundation. org/reports/disinformation-fake-news-and-influence-campaigns-on-twitter/
- [29] Aman Framewala, Aum Patil, and Faruk Kazi. 2020. Shapley based Interpretable Semi-supervised Model for Detecting Similarity Index of Social Media Campaigns. In 2020 International Conference on Computational Science and Computational Intelligence (CSCI). 270–274. https: //doi.org/10.1109/CSCI51800.2020.00052
- [30] Stéphane Ganassali. 2008. The Influence of the Design of Web Survey Questionnaires on the Quality of Responses. Survey Research Methods 2, 1 (mar 2008), 21–32. https://doi.org/10.18148/srm/2008.v2i1.598
- [31] Bhavya Ghai, Q. Vera Liao, Yunfeng Zhang, Rachel Bellamy, and Klaus Mueller. 2021. Explainable Active Learning (XAL): Toward AI Explanations as Interfaces for Machine Teachers. Proc. ACM Hum.-Comput. Interact. 4, CSCW3, Article 235 (Jan. 2021), 28 pages. https://doi.org/10.1145/3432934
- [32] Mark E. Glickman, Sowmya R. Rao, and Mark R. Schultz. 2014. False discovery rate control is a recommended alternative to Bonferroni-type adjustments in health studies. *Journal of Clinical Epidemiology* 67, 8 (2014), 850–857. https://doi.org/10.1016/j.jclinepi.2014.03.012
- [33] David Gunning, Mark Stefik, Jaesik Choi, Timothy Miller, Simone Stumpf, and Guang-Zhong Yang. 2019. XAI—Explainable artificial intelligence. Science Robotics 4, 37 (2019), eaay7120. https://doi.org/10.1126/scirobotics.aay7120 arXiv:https://www.science.org/doi/pdf/10.1126/scirobotics.aay7120
- [34] Ronan Hamon, Henrik Junklewitz, Gianclaudio Malgieri, Paul De Hert, Laurent Beslay, and Ignacio Sanchez. 2021. Impossible Explanations? Beyond Explainable AI in the GDPR from a COVID-19 Use Case Scenario. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (Virtual Event, Canada) (FAccT '21). Association for Computing Machinery, New York, NY, USA, 549–559. https://doi.org/10.1145/ 3442188.3445917
- [35] Casper Hansen, Christian Hansen, Stephen Alstrup, Jakob Grue Simonsen, and Christina Lioma. 2019. Neural Check-Worthiness Ranking with Weak Supervision: Finding Sentences for Fact-Checking. In Companion Proceedings of The 2019 World Wide Web Conference (San Francisco, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 994–1000. https://doi.org/10.1145/3308560.3316736

Gionnieve Lim and Simon T. Perrault

- [36] McKenzie Himelein-Wachowiak, Salvatore Giorgi, Amanda Devoto, Muhammad Rahman, Lyle Ungar, H Andrew Schwartz, David H Epstein, Lorenzo Leggio, and Brenda Curtis. 2021. Bots and Misinformation Spread on Social Media: Implications for COVID-19. J Med Internet Res 23, 5 (May 2021). https://doi.org/10.2196/26933
- [37] Reuters Institute. 2022. Digital News Report 2022. Retrieved June 1, 2023 from https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2022
- [38] Md Rafiqul Islam, Shaowu Liu, Xianzhi Wang, and Guandong Xu. 2020. Deep learning for misinformation detection on online social networks: a survey and new perspectives. Social Network Analysis and Mining 10, 82 (Sept. 2020). https://doi.org/10.1007/s13278-020-00696-x
- [39] Jasser Jasser. 2019. Dynamics of Misinformation Cascades. In Companion Proceedings of The 2019 World Wide Web Conference (San Francisco, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 33–36. https://doi.org/10.1145/3308560.3314194
- [40] Caio Libanio Melo Jeronimo, Leandro Balby Marinho, Claudio E. C. Campelo, Adriano Veloso, and Allan Sales da Costa Melo. 2019. Fake News Classification Based on Subjective Language. In Proceedings of the 21st International Conference on Information Integration and Web-Based Applications & Services (Munich, Germany) (iiWAS2019). Association for Computing Machinery, New York, NY, USA, 15–24. https://doi.org/10.1145/3366030.3366039
- [41] Tanveer Khan, Antonis Michalas, and Adnan Akhunzada. 2021. Fake news outbreak 2021: Can we stop the viral spread? Journal of Network and Computer Applications 190 (2021), 103112. https://doi.org/10.1016/j.jnca.2021.103112
- [42] Jan Kirchner and Christian Reuter. 2020. Countering Fake News: A Comparison of Possible Solutions Regarding User Acceptance and Effectiveness. Proc. ACM Hum.-Comput. Interact. 4, CSCW2, Article 140 (Oct. 2020), 27 pages. https://doi.org/10.1145/3415211
- [43] Vivian Lai and Chenhao Tan. 2019. On Human Predictions with Explanations and Predictions of Machine Learning Models: A Case Study on Deception Detection. In Proceedings of the Conference on Fairness, Accountability, and Transparency (Atlanta, GA, USA) (FAT* '19). Association for Computing Machinery, New York, NY, USA, 29–38. https://doi.org/10.1145/3287560.3287590
- [44] Markus Langer, Daniel Oster, Timo Speith, Holger Hermanns, Lena Kästner, Eva Schmidt, Andreas Sesing, and Kevin Baum. 2021. What do we want from Explainable Artificial Intelligence (XAI)? – A stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research. *Artificial Intelligence* 296 (2021), 103473. https://doi.org/10.1016/j.artint.2021.103473
- [45] Gionnieve Lim and Simon T. Perrault. 2022. Explanation Preferences in XAI Fact-Checkers. Proceedings of 20th European Conference on Computer-Supported Cooperative Work (2022). https://doi.org/10.48340/ecscw2022_p02
- [46] Scott M. Lundberg and Su-In Lee. 2017. A Unified Approach to Interpreting Model Predictions. In Proceedings of the 31st International Conference on Neural Information Processing Systems (Long Beach, California, USA) (NIPS'17). Curran Associates Inc., Red Hook, NY, USA, 4768–4777. https://doi.org/10.5555/3295222.3295230
- [47] Harrison Mantas. 2021. Twitter finally turns to the experts on fact-checking. Retrieved June 1, 2023 from https://www.poynter.org/fact-checking/2021/twitter-finally-turns-to-the-experts-on-fact-checking
- [48] Julian Marx, Felix Brünker, Milad Mirbabaie, and Eric Hochstrate. 2020. 'Conspiracy Machines' The Role of Social Bots during the COVID-19 'Infodemic'. ACIS 2020 Proceedings 82 (2020). https://aisel.aisnet.org/acis2020/82
- [49] Yelena Mejova and Kyriaki Kalimeri. 2020. COVID-19 on Facebook Ads: Competing Agendas around a Public Health Crisis. In Proceedings of the 3rd ACM SIGCAS Conference on Computing and Sustainable Societies (Ecuador) (COMPASS '20). Association for Computing Machinery, New York, NY, USA, 22–31. https://doi.org/10.1145/3378393.3402241
- [50] An T. Nguyen, Aditya Kharosekar, Saumyaa Krishnan, Siddhesh Krishnan, Elizabeth Tate, Byron C. Wallace, and Matthew Lease. 2018. Believe It or Not: Designing a Human-AI Partnership for Mixed-Initiative Fact-Checking. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology* (Berlin, Germany) (UIST '18). Association for Computing Machinery, New York, NY, USA, 189–199. https: //doi.org/10.1145/3242587.3242666
- [51] Mahsan Nourani, Samia Kabir, Sina Mohseni, and Eric D. Ragan. 2019. The Effects of Meaningful and Meaningless Explanations on Trust and Perceived System Accuracy in Intelligent Systems. Proceedings of the AAAI Conference on Human Computation and Crowdsourcing 7, 1 (Oct. 2019), 97–105. https://doi.org/10.1609/hcomp.v7i1.5284
- [52] Gordon Pennycook, Adam Bear, Evan T. Collins, and David G. Rand. 2019. The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings. *Management Science* 66, 11 (2019), 4944–4957. https://doi.org/10.1287/ mnsc.2019.3478
- [53] Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2017. Where the Truth Lies: Explaining the Credibility of Emerging Claims on the Web and Social Media. In Proceedings of the 26th International Conference on World Wide Web Companion (Perth, Australia) (WWW '17 Companion). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1003–1012. https://doi.org/10.1145/3041021.3055133
- [54] Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2018. CredEye: A Credibility Lens for Analyzing and Explaining Misinformation. In Companion Proceedings of the The Web Conference 2018 (Lyon, France) (WWW '18). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 155–158. https://doi.org/10.1145/3184558.3186967
- [55] Facebook Journalism Project. 2021. How Facebook's third-party fact-checking program works. Retrieved June 1, 2023 from https://www.facebook. com/journalismproject/programs/third-party-fact-checking/how-it-works
- [56] Miriam Redi, Besnik Fetahu, Jonathan Morgan, and Dario Taraborelli. 2019. Citation Needed: A Taxonomy and Algorithmic Assessment of Wikipedia's Verifiability. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 1567–1578. https://doi.org/10.1145/3308558.3313618

Behavioral Effects and User Preferences on XAI Fact-Checkers

- [57] Julio C. S. Reis, André Correia, Fabrício Murai, Adriano Veloso, and Fabrício Benevenuto. 2019. Explainable Machine Learning for Fake News Detection. In Proceedings of the 10th ACM Conference on Web Science (Boston, Massachusetts, USA) (WebSci '19). Association for Computing Machinery, New York, NY, USA, 17–26. https://doi.org/10.1145/3292522.3326027
- [58] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (San Francisco, California, USA) (KDD '16). Association for Computing Machinery, New York, NY, USA, 1135–1144. https://doi.org/10.1145/2939672.2939778
- [59] Twitter Safety. 2020. COVID-19: Our approach to misleading vaccine information. Retrieved June 1, 2023 from https://blog.twitter.com/en_us/ topics/company/2020/covid19-vaccine
- [60] Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. 2019. Combating Fake News: A Survey on Identification and Mitigation Techniques. ACM Trans. Intell. Syst. Technol. 10, 3, Article 21 (April 2019), 42 pages. https://doi.org/10.1145/3305260
- [61] Ben Shneiderman. 2021. Human-Centered AI. Issues in Science and Technology 37, 2 (2021), 56-61. https://issues.org/human-centered-ai/
- [62] Kai Shu, Limeng Cui, Suhang Wang, Dongwon Lee, and Huan Liu. 2019. DEFEND: Explainable Fake News Detection. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Anchorage, AK, USA) (KDD '19). Association for Computing Machinery, New York, NY, USA, 395–405. https://doi.org/10.1145/3292500.3330935
- [63] Snopes. 2021. Fact Check Ratings. Retrieved June 1, 2023 from https://www.snopes.com/fact-check-ratings/
- [64] Tian Tian, Yudong Liu, Xiaoyu Yang, Yuefei Lyu, Xi Zhang, and Binxing Fang. 2020. QSAN: A Quantum-Probability Based Signed Attention Network for Explainable False Information Detection. Association for Computing Machinery, New York, NY, USA, 1445–1454. https://doi.org/10.1145/3340531. 3411890
- [65] TikTok. 2020. An update on our virtual Transparency & Accountability Center experience. Retrieved June 1, 2023 from https://newsroom.tiktok. com/en-us/an-update-on-our-virtual-transparency-and-accountability-center-experience
- [66] Jasper van der Waa, Elisabeth Nieuwburg, Anita Cremers, and Mark Neerincx. 2021. Evaluating XAI: A comparison of rule-based and example-based explanations. Artificial Intelligence 291 (2021), 103404. https://doi.org/10.1016/j.artint.2020.103404
- [67] Luis Vargas, Patrick Emami, and Patrick Traynor. 2020. On the Detection of Disinformation Campaign Activity with Network Analysis. In Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop (Virtual Event, USA) (CCSW'20). Association for Computing Machinery, New York, NY, USA, 133–146. https://doi.org/10.1145/3411495.3421363
- [68] Bairong Wang and Jun Zhuang. 2018. Rumor response, debunking response, and decision makings of misinformed Twitter users during disasters. Natural Hazards 93 (Sept. 2018), 1145–1162. https://doi.org/10.1007/s11069-018-3344-6
- [69] Patrick Wang, Rafael Angarita, and Ilaria Renna. 2018. Is This the Era of Misinformation yet: Combining Social Bots and Fake News to Deceive the Masses. In *Companion Proceedings of the The Web Conference 2018* (Lyon, France) (WWW '18). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1557–1561. https://doi.org/10.1145/3184558.3191610
- [70] Xinru Wang and Ming Yin. 2021. Are Explanations Helpful? A Comparative Study of the Effects of Explanations in AI-Assisted Decision-Making. In 26th International Conference on Intelligent User Interfaces (College Station, TX, USA) (IUI '21). Association for Computing Machinery, New York, NY, USA, 318–328. https://doi.org/10.1145/3397481.3450650
- [71] Yongyue Wang, Chunhe Xia, Chengxiang Si, Beitong Yao, and Tianbo Wang. 2020. Robust Reasoning Over Heterogeneous Textual Information for Fact Verification. *IEEE Access* 8 (2020), 157140–157150. https://doi.org/10.1109/ACCESS.2020.3019586
- [72] Katharina Weitz, Dominik Schiller, Ruben Schlagowski, Tobias Huber, and Elisabeth André. 2019. "Do You Trust Me?": Increasing User-Trust by Integrating Virtual Agents in Explainable AI Interaction Design. In Proceedings of the 19th ACM International Conference on Intelligent Virtual Agents (Paris, France) (IVA '19). Association for Computing Machinery, New York, NY, USA, 7–9. https://doi.org/10.1145/3308532.3329441
- [73] Jutta Williams and Rumman Chowdhury. 2021. Introducing our Responsible Machine Learning Initiative. Retrieved June 1, 2023 from https: //blog.twitter.com/en_us/topics/company/2021/introducing-responsible-machine-learning-initiative
- [74] Stephan Winter and Nicole C. Krämer. 2014. A question of credibility Effects of source cues and recommendations on information selection on news sites and blogs. Communications 39, 4 (2014), 435–456. https://doi.org/doi:10.1515/commun-2014-0020
- [75] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The Aligned Rank Transform for Nonparametric Factorial Analyses Using Only Anova Procedures. Association for Computing Machinery, New York, NY, USA, 143–146. https://doi.org/10.1145/1978942.1978963
- [76] Christine T. Wolf and Kathryn E. Ringland. 2020. Designing Accessible, Explainable AI (XAI) Experiences. SIGACCESS Access. Comput. 125, Article 6 (March 2020), 1 pages. https://doi.org/10.1145/3386296.3386302
- [77] Liang Wu, Fred Morstatter, Kathleen M. Carley, and Huan Liu. 2019. Misinformation in Social Media: Definition, Manipulation, and Detection. SIGKDD Explor. Newsl. 21, 2 (Nov. 2019), 80–90. https://doi.org/10.1145/3373464.3373475
- [78] Ming Yin, Jennifer Wortman Vaughan, and Hanna Wallach. 2019. Understanding the Effect of Accuracy on Trust in Machine Learning Models. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3290605.3300509
- [79] Xinyi Zhou and Reza Zafarani. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. ACM Comput. Surv. 53, 5, Article 109 (Sept. 2020), 40 pages. https://doi.org/10.1145/3395046



Fig. 6. The LIME explanation.

A XAI INTERFACES

Examples of the five XAI interfaces for the LIME (Figure 6), SHAP (Figure 7), Attention (Figure 8), Comments (Figure 9), and Evidence (Figure 10) explanations.

Behavioral Effects and User Preferences on XAI Fact-Checkers



Red and Blue Texts

The texts are features from the news content used to make the prediction. The "." symbol is used to shorten texts that are too long.

Red and Blue Arrows

The colour and length of the arrows represent the extent in which text features increase (red) or decrease (blue) the prediction.

Output and Base Values

The output value is the prediction made given the text features, while the base value is the baseline prediction made without them.

An output greater than the base value represents a higher prediction that the news is false, while a smaller output represents a lower prediction.

Fig. 7. The SHAP explanation.



Fig. 8. The Attention explanation.

Behavioral Effects and User Preferences on XAI Fact-Checkers

Fact-Checker Report						
Prediction made using Social Media Comments						
Top 3 Comments						
You mean the guy who also had his fist up in solidarity with the insurrection mob? Yeahth	is					
checks out. 😪	0.517					
There wasn't enough "treason" in the legislati whet his whistle.	on to 0.284					
He looks like a good guy. Let's give him the b of the doubt. 😂	enefit 0.210					

How to understand this?

Stance-Coloured Boxes

Green boxes are comments that support the claim of the news content, red boxes are against the claim, and yellow boxes are neutral.

Informativity Scores

The higher the score, the more informative and useful the comment. Comments are ranked in descending order of informativity.

Prediction

The 'true' or 'false' prediction is made based on the combined stance and informativity scores of the comments.

Fig. 9. The Comments explanation.

Fact-Checker Report

Prediction made using Evidence from Other News Sources

Top 3 Evidence

Senator Josh Hawley (R-MO) took to social media on Thursday to explain why he was the only senator to vote against a bill aimed at... 0.853

Hawley has distinguished himself among Senate Republicans by voting against every one of President Joe Biden's cabinet nominees... 0.748

Sen. Josh Hawley (R-Mo.) was the only senator to vote Thursday against legislation seeking to combat anti-Asian hate crimes, ... 0.322

How to understand this?

Stance-Coloured Boxes

Green boxes are evidence that support the claim of the news content and red boxes are evidence that do not.

Reliability Scores

The higher the score, the more reliable and trustworthy the evidence. Evidence are ranked in descending order of reliability.

Prediction

The 'true' or 'false' prediction is made based on the combined stance and reliability scores of the evidence.

Fig. 10. The Evidence explanation.