



Efficiency and Effectiveness of LLM-Based Summarization of Evidence in Crowdsourced Fact-Checking

Kevin Roitero
University of Udine
Udine, Italy
kevin.roitero@uniud.it

Dustin Wright
University of Copenhagen
Copenhagen, Denmark
dw@di.ku.dk

Michael Soprano
University of Udine
Udine, Italy
michael.soprano@uniud.it

Isabelle Augenstein
University of Copenhagen
Copenhagen, Denmark
augenstein@di.ku.dk

Stefano Mizzaro
University of Udine
Udine, Italy
stefano.mizzaro@uniud.it

Abstract

Evaluating the truthfulness of online content is critical for combating misinformation. This study examines the efficiency and effectiveness of crowdsourced truthfulness assessments through a comparative analysis of two approaches: one involving full-length webpages as evidence for each claim, and another using summaries for each evidence document generated with an LLM. Using an A/B testing setting, we engage a diverse pool of participants tasked with evaluating the truthfulness of statements under these conditions.

Our analysis explores both the quality of assessments and the behavioral patterns of participants. The results reveal that relying on summarized evidence offers comparable accuracy and error metrics to the standard modality while significantly improving efficiency. Workers in the *Summary* setting complete a significantly higher number of assessments, reducing task duration and costs. Additionally, the *Summary* modality maximizes internal agreement and maintains consistent reliance on and perceived usefulness of evidence, demonstrating its potential to streamline large-scale truthfulness evaluations.

CCS Concepts

- **Computing methodologies** → **Natural language processing;**
- **Information systems** → **Crowdsourcing.**

Keywords

Truthfulness Assessment, Summarization, Large Language Models

ACM Reference Format:

Kevin Roitero, Dustin Wright, Michael Soprano, Isabelle Augenstein, and Stefano Mizzaro. 2025. Efficiency and Effectiveness of LLM-Based Summarization of Evidence in Crowdsourced Fact-Checking. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3726302.3729960>



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGIR '25, Padua, Italy*

© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1592-1/2025/07
<https://doi.org/10.1145/3726302.3729960>

1 Introduction

The proliferation of online misinformation has made truthfulness assessment a crucial task in the digital era. Fact-checking efforts, traditionally conducted by experts, have faced challenges in scalability and timeliness as the volume of information continues to grow [5]. Recent developments, such as X and Meta scaling down their reliance on paid fact-checkers,¹ further highlight the urgency of scalable and efficient truthfulness assessment methods [7]. As social media platforms reduce investments in professional fact-checking, alternative approaches become increasingly vital for addressing the growing volume of misinformation. To address this, crowdsourcing has emerged as a promising alternative, leveraging the collective efforts of non-expert workers to evaluate truthfulness efficiently [35, 45]. However, the success of crowdsourced truthfulness assessments relies heavily on the format, presentation, and quality of evidence provided to participants. The choice of evidence itself, along with how it is presented, directly influences workers' ability to make accurate and reliable judgments.

Recent advancements in Large Language Models (LLMs) have opened new possibilities for improving the efficiency of truthfulness assessments by generating summarized evidence from lengthy documents. These models can distill essential information into concise formats, potentially reducing cognitive load and speeding up decision-making for crowd workers. Summarization also allows workers to process more evidence in less time, which is particularly valuable in large-scale annotation tasks. However, the adoption of LLM-generated summaries introduces its own set of challenges, including the risk of omitting critical details, introducing biases, or oversimplifying nuanced arguments. While summarized evidence could offer advantages in terms of efficiency and scalability, it may also affect the quality of truthfulness assessments. Workers might rely too heavily on the condensed summaries, overlooking important context available in full-length documents. Additionally, inconsistencies or factual errors introduced during the summarization process could inadvertently mislead workers, potentially impacting the reliability of their assessments. These trade-offs necessitate a thorough investigation into both the benefits and limitations of using LLM-generated summaries in crowdsourced truthfulness evaluations.

¹<https://www.nytimes.com/2025/01/07/business/meta-community-notes-x.html>.

In this work, we focus on comparing two approaches: a *Standard* modality, in which participants evaluate full-length webpages, and a *Summary* modality, in which participants assess summarized evidence generated by a state-of-the-art LLM, Meta-Llama-3-8B-Instruct [15]. The summarized format aims to reduce cognitive load and improve worker efficiency by condensing essential information while preserving factual accuracy. This presentation allows participants to process more information in less time, increasing overall throughput. Such improvements not only offer productivity gains but also translate into significant cost savings for large-scale annotation tasks, making summarization particularly valuable. However, the effectiveness and trade-offs of these two modalities remain an open question, which we address through a detailed empirical comparison using an A/B testing framework. This experimental design allows us to systematically compare the two modalities under controlled conditions, providing insights into how evidence presentation affects truthfulness evaluations. To guide our analysis, we address the following Research Questions (RQs):

- RQ1: Effectiveness.** Does the use of summarized evidence achieve comparable accuracy and error metrics to the standard approach of collecting judgments with full-length webpages?
- RQ2: Efficiency.** How does the use of summarized evidence impact the time required for participants to complete truthfulness assessments compared to full-length evidence? To what extent can summarization reduce the costs associated with large-scale truthfulness assessment tasks?
- RQ3: Behavioral Insights.** How does the presentation of evidence influence participants' behavioral patterns, such as reliance on evidence?

The remainder of this paper is structured as follows: Section 2 reviews the background and related work. Section 3 describes our experimental design, including the preparation of data, generation of summarized evidence, the structure of crowdsourcing tasks, and provides demographic statistics. The results are then presented across three key dimensions: effectiveness (Section 4), efficiency (Section 5), and worker behavior (Section 6). Finally, Section 7 discusses the broader implications of our findings, acknowledges the limitations of the study, and suggests directions for future work.

2 Background and Related Work

Multiple studies have explored the application of crowdsourcing for misinformation detection, demonstrating its potential as a scalable alternative to expert fact-checking [29, 40]. La Barbera et al. [28] highlighted the influence of worker bias in assessing statements, an issue further explored by Roitero et al. [45], who observed similar agreement levels across various truthfulness scales. Expanding on this, Soprano et al. [50] and Liu et al. [32] introduced multidimensional truthfulness scales, showing that crowd judgments are reliable across multiple truthfulness dimensions, each capturing different aspects of truthfulness. The comparative performance of crowd workers and experts has also been examined. For example, Zhao and Naaman [61] found that while the crowd performed on par with experts in terms of accuracy, objectivity, and clarity, their assessments often depended on existing professional knowledge. Similarly, Saeed et al. [46] analyzed social network data and demonstrated that crowd evaluations relied on different evidence sources

compared to experts, with better scalability and efficiency. Allen et al. [2] further highlighted the scalability of crowd-based methods, noting that political alignment influenced crowd judgments. Additionally, linguistic [55] and psychological [41] peculiarities in misinformation were explored, emphasizing the diverse factors affecting the quality of crowdsourced truthfulness evaluations.

While crowdsourcing holds significant potential, it is not without challenges [17, 25]. Draws et al. [13] identified a general tendency among workers to overestimate truthfulness, often driven by systematic cognitive biases [14, 49]. Bias-aware aggregation methods have been proposed to mitigate these distortions, such as the confirmation bias addressed by Gemalmaz and Yin [22]. Additionally, Nguyen et al. [38] presented FactCatch, a human-in-the-loop system designed to minimize effort while guiding users in fact-checking tasks. Further analyses have explored how biases can be identified and corrected to improve the quality of crowdsourced assessments. For example, Han et al. [23] investigated task abandonment patterns, while Li et al. [30] proposed methods to enhance labeling reliability by leveraging correlations between labels and neighboring instances. These efforts underscore the importance of addressing cognitive and procedural challenges to maximize the reliability of crowdsourced truthfulness assessments.

Although this study focuses on crowdsourcing, advances in automatic and hybrid fact-checking methods have contributed significantly to the field. Fully automated techniques, mostly based on machine learning models, have been widely explored for misinformation detection [1, 4, 5, 10, 11, 24, 34]. These methods have been evaluated using public datasets and benchmarks, including those provided by initiatives like the CLEF CheckThat! Lab [36, 37] as well as other projects such as MultiFC [6], FEVER [52, 53], AVeriTeC [47] and Factcheck-Bench [56]. Comparative analyses also have been developed to study variability in model performance across datasets [24, 43, 48, 51]. Recent research has also focused on generating human-readable explanations [3, 8, 54, 59] and incorporating complex interactions within machine learning pipelines [31]. Hybrid approaches that integrate human input with machine learning have gained attention as well [12, 42, 58]. Lastly, there is research on how fact checkers use tools in practice [57].

In this paper, we leverage automatic signals generated by a state-of-the-art LLM, Meta-Llama-3-8B-Instruct [15], to enhance the efficiency and scalability of a crowd-based approach for truthfulness assessment. By combining LLM-generated summaries with crowdsourced evaluations, we aim to investigate the interplay between automated summarization and human judgment in achieving reliable and cost-effective misinformation detection.

3 Methodology

In this section, we detail the methodology employed in our experiment. To promote transparency and reproducibility, all data collected during the experiment are made publicly available to the research community.²

3.1 Data

This study uses the same dataset and adopts a similar experimental design to that reported by La Barbera et al. [27], with modifications

²<https://doi.org/10.17605/OSF.IO/BNFXM>.

Table 1: Example of an LLM-generated summary given the query ‘Opposition to having a fully elected Chicago Board of Education is in the “super minority.”’.

Evidence	Summarized Evidence
For the first time in Chicago’s history, voters would get a say in who runs the city’s school board under two competing proposals now before the state Legislature. One calls for a fully elected Board of Education, and the other for a “hybrid” model splitting the school board into some elected members with the majority still appointed by the mayor. [...]	<ul style="list-style-type: none"> - The document discusses the possibility of having a fully elected Chicago Board of Education, with a focus on the overwhelming support for this idea among Chicagoans, as shown through various polls and referendums. - The current system of political appointment by the mayor has been in place for 150 years, but there is no consensus among researchers on which form of governance fosters better student performance or fiscal management. - Over 90% of school districts nationwide have elected boards, and overwhelming majorities of Chicagoans have long favored a switch to an elected board. [...]

tailored to meet the specific objectives and requirements of our research. To ensure that this paper is self-contained and understandable independently, we report both the details of the dataset and the crowdsourcing task employed. The dataset consists of a subset of statements sourced from the PolitiFact website,³ as in previous research [13, 27, 50]. PolitiFact has been an active platform for fact-checking statements made by U.S. politicians, political organizations, public figures, and social media posts since 2007, accumulating over 24,000 fact-checks with regular updates.

PolitiFact statements are categorized by expert judges using a six-level truthfulness scale, consisting of the following levels: Pants On Fire, False, Mostly False, Half True, Mostly True, and True. In this study, we adopt the same six-level scale to maintain consistency and enable direct comparisons between expert fact-checkers and crowdsourced assessments. As in previous work, we select a subset of 120 statements (20 for each ground truth category) from PolitiFact. Additionally, we ensure that each statement is supported by at least 10 webpages retrieved as evidence using the Bing Web Search API.⁴ These statements form the basis of the crowdsourcing experiments described in the subsequent sections.

3.2 Evidence Summarization

We process the evidence related to each statement to create two distinct versions for evaluation: the original, unaltered full-length web page and a more concise summarized version, obtained by leveraging a LLM. To achieve the summarization version, we employ the capabilities of the Meta-Llama-3-8B-Instruct model, a state-of-the-art model known for its robust performance in natural language understanding and generation tasks [15].⁵

The process begins with a carefully crafted prompt provided to the model. This prompt is specifically designed to guide the model in extracting and summarizing the essential factual elements from

³<https://www.politifact.com/>

⁴<https://www.microsoft.com/en-us/bing/apis/bing-web-search-api>

⁵<https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

System Prompt:

You are a helpful, respectful, and honest assistant. Your job is to provide a summary of the provided documents based on a user query. Your summaries should be as accurate as possible and should not include any harmful, unethical, racist, sexist, toxic, dangerous, or illegal content. Please ensure that your responses are socially unbiased and positive in nature.

User Prompt:

Here is a user query: {statement}. Please do your best to provide a concise list of points which summarize the following document. Please focus on points related to the query. Please output your response as a JSON object with a field "summary" that contains the list of strings. Please first describe what the document is. Then please output five or more strings summarizing the contents of the document. Here is the document: {document content}.

Figure 1: Evidence summarization prompt.

the content of a webpage, ensuring that the generated summary maintains consistency with the original information in terms of factuality and stance. After several rounds of trial and error, the exact wording of the prompt we use is given in Figure 1.

Using this prompt, the model generates a summary that distills the webpage’s content into a shorter form, focusing on the most critical information needed to perform the fact-checking task. An example of the summarized evidence is shown in Table 1.

3.3 Crowdsourcing Tasks

We adopt the crowdsourcing task as detailed in La Barbera et al. [27], modifying it to collect responses for our study while preserving the core structure of the original design. This design, which was originally developed and validated in earlier research [13, 27, 50], allows us to collect truthfulness assessments at scale using crowdsourcing.

We employ a task design featuring sequential data collection, where participants evaluate statements one at a time for truthfulness. We recruit participants through the Prolific platform and apply a filter to include only workers from the U.S. Compensation reflects fair market practices and is based on careful estimations of task completion times, ensuring that payments are both ethical and practical. For the *Standard* modality, we pay each worker £2, while for the *Summary* modality, we set compensation at £1.80. Each worker evaluates a total of 8 statements (one for each ground truth category and two additional gold-standard questions), and 5 distinct workers assess each statement. We recruit 100 workers for each modality, resulting in a total cost of approximately £500, including Prolific’s fees.

After filling a demographic questionnaire, the participants perform our task using a structured interface where they first encounter the statement to be assessed, followed by a set of evidence. Each worker sees the title of the web page along with a snippet, similar to the presentation layout of a search engine. Then, depending on the task layout, if the worker clicks on the link they see either the full webpage or its summary. Subsequent to evidence review and selection, participants assess the statement. We ask

Table 2: Demographic distribution across modalities.

Consideration	Standard (%)	Summary (%)
Democrat	40	45
Independent	17	23
Republican	40	31
Something else	3	1
Political Views	Standard (%)	Summary (%)
Conservative	25	23
Liberal	24	26
Moderate	23	25
Very conservative	11	8
Very liberal	17	18

workers to first evaluate the overall truthfulness of the statement using the same six-level scale as the one used by experts. Then, we ask the workers to answer two questions regarding how useful they found the evidence in determining the truthfulness of the statement: “How useful was the evidence?” and “Did you use the evidence to fact-check the statement, or did you have a clear opinion before reviewing the evidence?”. For both questions, we collect responses on a five-level Likert scale. The scale for the first question includes the following options: “Not Useful At All”, “Slightly Useful”, “Moderately Useful”, “Very Useful”, and “Extremely Useful”. The scale for the second question includes: “I did not consider the evidence at all before forming my opinion”, “I considered the evidence slightly before forming my opinion”, “I considered the evidence as much as my initial opinion before forming my final opinion”, “I considered the evidence more than my initial opinion before forming my final opinion”, and “I fully based my opinion on the evidence after reviewing it”. To ensure high-quality assessments from workers, we implement the following checks: a minimum time requirement of 3 seconds per statement and correct evaluations of two gold-standard questions, one clearly true and the other clearly false.

3.4 Descriptive Statistics

The answers to the demographic questionnaire enable us to assess the demographics of participants across the two experimental conditions. Table 2 presents the distribution of political considerations and views for each condition. The results indicate a balanced representation across both modalities, with comparable percentages for both political affiliation and ideological orientation.

We also measure abandonment [23] to evaluate participant engagement. Initial abandonment rates (49% for *Standard*, 38% for *Summary*) and mid-task abandonment rates (10% for *Standard*, 13% for *Summary*) are comparable, with no statistically significant differences observed.

4 RQ1: Effectiveness

4.1 Agreement with Experts

We inspect the agreement between workers and experts, shown in Figure 2, where we report results for both individual and aggregated assessments. In all figures, the six PolitiFact levels are abbreviated with their initials. We discuss both the individual and

the aggregated scores for the *Standard* and *Summary* modalities. In terms of individual scores, in the standard modality, we observe an accuracy of 0.31, a MSE of 3.06, and a MAE of 1.27. The summary modality reports a similar accuracy (0.31) but exhibits a slightly higher MSE (3.27) and MAE (1.30). This similarity in the metric scores suggests a comparable quality level between the standard and summary modalities, indicating that the summarized content does not impact workers’ effectiveness. This finding is nontrivial, given that the content presented in the summary version of the task is substantially condensed. We further examined whether the observed differences in MAE and MSE between the modalities are statistically significant. Using bootstrapped confidence intervals⁶ with 10,000 resamples, we estimated the range of plausible values for the difference in these metrics. For MAE, the mean difference was -0.028 with a 95% confidence interval ranging from -0.167 to 0.112 . Similarly, for MSE, the mean difference was -0.212 with a 95% confidence interval spanning from -0.767 to 0.340 . As both intervals include zero, we conclude that there is no statistically significant difference between the standard and summary modalities in terms of MAE and MSE.

We now turn to the aggregated scores, which are still shown in Figure 2. The scores are aggregated using three distinct statistical methods: mean, mode (i.e., majority voting), and median. The mean offers a general view of central tendency and has proven to lead to higher agreement with expert assessments [27, 50], while the mode identifies the most frequently occurring outcome, and the median provides an additional robust measure of central tendency that is less sensitive to extreme values. In the standard modality, mean accuracy is slightly higher (0.33) compared to the summary modality (0.31), though the difference is minimal. Similarly, the standard modality shows marginally better error metrics, with MSE (1.65 vs. 1.77) and MAE (0.95 vs. 1.00) indicating lower dispersion and deviation from the ground truth. Mode aggregation reflects a similar trend, with accuracy at 0.42 for the standard modality and 0.38 for the summary modality, alongside smaller errors (MSE: 1.81 vs. 2.13, MAE: 0.91 vs. 0.98). Median aggregation further supports this finding, with the standard modality showing slightly better performance across metrics (accuracy: 0.40 vs. 0.38, MSE: 1.80 vs. 2.13, MAE: 0.92 vs. 0.93). While errors in the summary modality are slightly more pronounced, both approaches achieve comparable accuracy levels. Bootstrap analysis confirms no statistically significant differences across accuracy and error measures. Overall, both modalities show similar effectiveness across several metrics, indicating that, despite the reduced volume of information provided to workers in the summary modality, it closely resembles the standard modality in terms of data quality. To test whether worker demographics influence the effectiveness of truthfulness assessments, we additionally run an Ordinary Least Squares (OLS) regression analysis [21]. The results show no statistically significant relationship between demographics and accuracy, indicating that these factors do not substantially impact assessment performance.

4.2 Agreement Among Workers

We employ Krippendorff’s α [26] to assess the internal agreement among participants in the two task variations. Krippendorff’s α is

⁶<https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.bootstrap.html>.

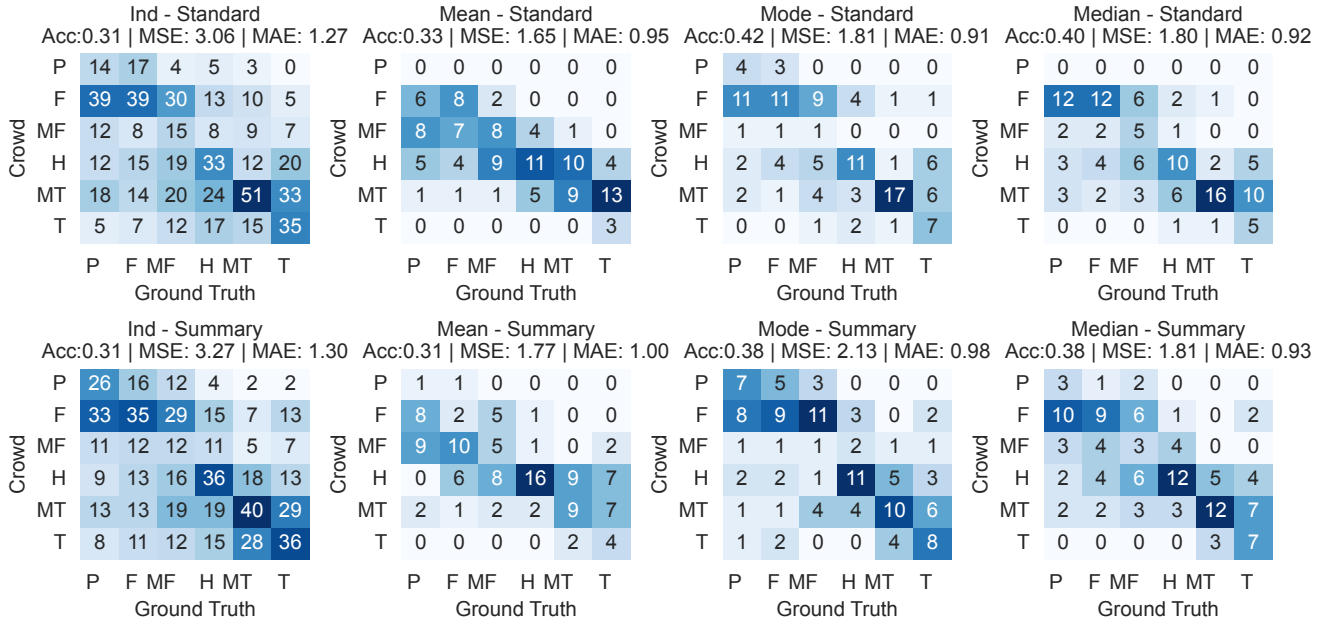


Figure 2: Agreement between workers and experts on individual and aggregated assessments.




Figure 3: Krippendorff's α score.

a robust statistical measure used to determine the consistency of ratings provided by different evaluators, and it is a very popular measure for crowdsourcing tasks [9].

Figure 3 shows the values of internal agreement across the two task modalities, both for the whole task (dotted lines) as well as for the single ground truth levels of the statements. For the *Standard* modality, the α values are generally lower than the *Summary* modality, reflecting overall less agreement among participants, with the only exception of the Pants On Fire truthfulness level. This suggests that when participants are presented with full-length content, there may be greater diversity in how the information is interpreted, possibly due to the complexity and volume of the provided information. Furthermore, the lowest α values appear in the categories in the middle of the scale, indicating that the claims with no extreme truthfulness score are also the most challenging for participants to consistently evaluate. Conversely, the *Summary* modality shows higher α values across all categories except Pants On Fire and for the overall task. This difference is statistically significant at the $p < 0.05$ level, according to the Mann-Whitney U Test [33],

indicating a higher degree of consistency among workers' ratings. This improvement in internal agreement could be attributed to the condensed nature of the content, which might focus the workers' evaluations and reduce both the cognitive load as well as the ambiguity or variability in interpretation that more extensive evidence might introduce. Notably, the alpha values are particularly higher in categories like Mostly False and True, where the quality of the summarization may help clarify the key aspects of the statements.

Overall, these findings remark the impact of evidence summarization on the consistency of workers' assessments. The higher internal agreement in the summary modality suggests that summarization not only maintains the essential content needed for the truthfulness assessment but may also aid in aligning worker scores, leading to more uniform interpretations and assessments of the truthfulness of statements.


4.3 Error Analysis

To better understand the patterns of errors in truthfulness assessments, we conducted a failure analysis by categorizing worker assessments into three types: *correct*, *overestimation*, and *underestimation*. The results, shown in the first part of Table 3 indicate that the distribution of error types is similar across the two modalities. These comparable distributions suggest that the modality of evidence presentation does not significantly affect the overall patterns of worker errors.

A more fine-grained analysis, shown in the second part of Table 3, highlights the distribution of errors across modalities. In the *Standard* modality, *overestimation* emerges as the most frequent error type for categories such as Mostly False (62.2%). Conversely, *correct* assessments are most common in categories like Mostly True (50.0%) and Half True (48.5%), suggesting workers find it

Table 3: Error types across truthfulness levels.


Modality	Scope	Correct	Over	Under
<i>Standard</i>	Overall	187 (30.6%)	237 (38.8%)	176 (28.8%)
<i>Summary</i>	Overall	185 (30.2%)	232 (37.9%)	183 (29.9%)
<i>Standard</i>	Pants On Fire	24 (24.0%)	76 (76.0%)	–
	False	39 (37.1%)	44 (41.9%)	35 (21.0%)
	Mostly False	15 (18.3%)	51 (62.2%)	34 (19.5%)
	Half True	33 (48.5%)	41 (30.1%)	26 (20.6%)
	Mostly True	51 (50.0%)	15 (14.7%)	34 (33.3%)
<i>Summary</i>	True	35 (35.0%)	–	65 (65.0%)
	Pants On Fire	26 (26.0%)	74 (74.0%)	–
	False	35 (35.7%)	49 (50.0%)	16 (14.3%)
	Mostly False	12 (15.6%)	47 (61.0%)	41 (23.4%)
	Half True	36 (42.9%)	34 (40.5%)	13 (16.6%)
<i>Summary</i>	Mostly True	40 (50.6%)	28 (32.0%)	32 (17.4%)
	True	36 (36.0%)	–	64 (64.0%)

**Figure 4: Accuracy across sampling sizes.**

easier to correctly evaluate statements closer to the middle of the truthfulness spectrum. *Underestimation* occurs less frequently overall but is more pronounced for higher truthfulness categories. In the *Summary* modality, *overestimation* remains dominant in the same categories as in the *Standard* modality. However, *correct* assessments are slightly more evenly distributed, with the highest rates observed for *Mostly True* (50.6%) and *Half True* (42.9%). Overall, these findings suggest that both modalities exhibit similar patterns with respect to error types. We further analyze the mismatch magnitude for *overestimation* and *underestimation* errors across the *Standard* and *Summary* modalities and find similar error patterns. Both modalities exhibit comparable average mismatch values for *overestimation* (1.99) and slightly more negative values for *underestimation* in the *Summary* modality (-1.73 vs. -1.65). None of these differences are statistically significant, confirming the overall comparability of the two approaches.

4.4 Robustness to Fewer Judgments

To investigate the robustness of the agreement between worker and expert truthfulness evaluations, we conduct a bootstrap analysis [16]. This method involves resampling the collected assessments

**Figure 5: Time elapsed on statements for the two modalities.**

with replacement, enabling us to estimate the variability in effectiveness metrics across different sampling sizes. By applying this approach, we aim to understand how the number of workers per document impacts the overall accuracy and consistency of aggregated truthfulness assessments in both task modalities. For each sampling size, ranging from 1 to 4 workers per document (5 would include all workers), we perform 1,000 repetitions during which we compute the individual accuracy of the aggregated assessments compared to expert evaluations.

The results, shown in Figure 4, reveal a trend in which the accuracy of both modalities slightly improves, although without statistical significance, and the variance decreases as the number of workers per document increases. This pattern is consistent with expectations from crowdsourcing settings, where aggregating multiple assessments tends to mitigate individual biases and errors. Notably, the *Summary* modality achieves effectiveness levels comparable to the *Standard* modality across all sampling sizes. These indicate that summarized evidence can provide comparable results to full-length evidence in terms of both accuracy and consistency, and that this equivalence is achieved while potentially reducing the cognitive load on workers and the resources required for task completion.

5 RQ2: Efficiency

5.1 Time

In evaluating the practical aspects of crowdsourced truthfulness assessments, time efficiency emerges as a critical metric, as the more time a worker spends on the task, the higher the cost required for their judgment and, consequently, for the overall task. Figure 5 shows the time required for workers to provide assessments for both task modalities. As shown by inspecting the two series in the plot, the average time taken to complete assessments consistently decreases as the statement index increases (i.e., as a worker performs the task). This trend suggests that participants become more proficient at navigating and processing the tasks over time. Such a learning effect, also observed in previous studies [27, 50], suggests the adaptability of participants and their increasing efficiency in handling the assessments, whether they are dealing with full-length or summarized content.

When comparing the two series we see that the *Summary* modality consistently shows a reduced average time on each statement

index compared to the *Standard* modality, with the only exception being the second document, and the difference is statistically significant according to the Mann-Whitney U Test [33] at the $p < 0.05$ level. This supports the observation that time efficiency gained from presenting workers with summarized evidence is real. In practice, the reduced time spent in the *Summary* modality can lead to enhanced productivity and a lighter cognitive load for participants, who can perform more judgments in the same amount of time and are required to read much less text. Such improvements in efficiency could also translate into cost savings for researchers and practitioners implementing these misinformation assessments at scale, potentially enhancing the economic feasibility of deploying crowdsourced truthfulness evaluations. This can be seen by inspecting the cumulative series in Figure 5, which show the potential for cost savings in large-scale truthfulness assessments.

For instance, completing 600 assessments, equivalent to 120 documents, requires approximately 23.67 hours of total work and costing £171.60 at the U.S. minimum wage of \$7.25 per hour (approximately £6.08). In contrast, the *Summary* modality, reduces the required total work to approximately 20.50 hours, costing £148.63 under the same wage assumption. While these figures are based on the minimum wage, actual costs would scale proportionally if higher wages are offered. When extended to larger datasets requiring (let us say) 6,000 assessments, the *Standard* modality incurs costs of £1,716, compared to £1,486 for the *Summary* modality, resulting in higher savings. These findings emphasize the economic advantages of the *Summary* modality, which maintains comparable performance metrics while significantly reducing worker time and overall costs, making it particularly valuable for large-scale annotation efforts.

5.2 Effectiveness Comparison Under Equal Time Constraints

To evaluate the effectiveness of evidence summarization under equal time constraints, we compared the effectiveness of the two modalities across metrics. Additionally, we assessed the efficiency of each modality in terms of the number of assessments completed within the same time frame, highlighting the potential advantages of summarization for throughput. To ensure a fair comparison, we simulate equal time constraints for both modalities by aligning the total time spent on assessments for each document in the dataset. For every document in the *Standard* modality, we compute the total time workers spend on assessments. Then, we iteratively stack assessments from the corresponding *Summary* modality until their cumulative time matches or exceeds the total time recorded for the *Standard* modality. Then, we compute metrics separately for each document in both modalities. We use the *Standard* modality as the baseline and express the relative efficiency of the *Summary* modality as a judgment multiplier, calculated as the ratio of completed assessments in the *Summary* modality to those in the *Standard* modality.

Table 4 presents the comparative results. The *Standard* modality achieves a slightly higher accuracy (0.312) compared to the *Summary* modality (0.308), with minimal differences in error metrics. These findings indicate comparable performance in judgment quality between the two modalities, also considering that no statistical significance is detected. The primary advantage of the *Summary*

Table 4: Comparison under equal time constraints.

Metric	Standard Value	Summary Value
Accuracy	0.312	0.308
MSE	3.063	3.275
MAE	1.270	1.298
Judgment Multiplier	–	+15%

modality thus lies in its efficiency. On average, workers in the *Summary* setting completed about 15% more assessments than those in the *Standard* setting, offering significant efficiency gains and cost savings, as detailed earlier.

Previously, we used the *Standard* modality as the baseline, stacking judgments in the *Summary* modality to match the total time spent on assessments in the *Standard* batch. In this analysis, we reverse the approach by sampling judgments from the *Standard* modality to match the total time spent in the *Summary* batch, using the *Summary* modality as the baseline. We do that by using a simulation-based analysis. In each of the 1,000 simulations, we randomly sample assessments from the *Standard* modality to match the total time spent in the *Summary* modality. The results are shown in Figure 6. The average effectiveness scores achieved by the *Standard* modality across the simulations indicates that the *Standard* modality maintains a comparable level of accuracy even when constrained by the same amount of time. As in the previous case, no statistical significance emerges between the two modalities. These findings suggest that both modalities perform similarly in terms of absolute deviations from the ground truth, with no significant disparities in judgment quality.

6 RQ3: Behavioral Insights

6.1 Usefulness and Need for Evidence

In the context of crowdsourced truthfulness assessments, understanding the perceived usefulness of the evidence is important for understanding the quality and trustworthiness of the workers' provided assessments. This provides insight into how the evidence impacts the decision-making process and the reliance of participants on the provided information. Figure 7 shows the perceived usefulness and necessity of the evidence for the two modalities.

We start with inspecting the need for evidence as expressed by participants across different truthfulness levels. Data show a significant reliance on evidence, with the majority of responses across both *Standard* and *Summary* modalities suggesting that assessments are “fully based on evidence”. This strong dependency highlights the critical role that evidence plays in ensuring accurate and reliable assessments. Notably, in the summary modality, there is a slightly higher tendency for assessments to be fully based on evidence compared to the standard modality, suggesting that even summarized content sufficiently supports robust judgment formation. Less frequently, participants report that their assessments were “more than my initial opinion” or “considered initial opinion”, indicating that while evidence guides their evaluations, preconceived notions or initial impressions also play a role. Interestingly, the number of participants who report their assessments as “considered slightly” or “did not consider” the evidence is minimal,




Figure 6: Distributions of Accuracy, Mean Squared Error (MSE), and Mean Absolute Error (MAE) across simulations.




Figure 7: Need and perceived usefulness of the evidence in the task.

suggesting the minimal influence of unsupported personal bias in the evaluation process.

Turning to the usefulness of evidence, participants across both modalities rate the majority of the evidence as being “extremely useful” or “very useful”. This high level of perceived usefulness is consistent across all truthfulness categories, remarking the importance of quality evidence in the assessment task. In the *Standard* modality, there is a notable pattern where evidence is frequently considered at least “moderately useful”, which aligns with the reduced amount of information presented to the workers, yet keeping its impact on their assessments. The less frequent ratings of “slightly useful” or “not useful at all” are significantly lower, which indicates that the evidence provided is overall useful to workers. This holds as well for the *Summary* modality: also summarized information appears to be adequately informative to workers.

Overall, these findings suggest the fundamental role of evidence in shaping the outcomes of crowdsourced truthfulness assessments. The high reliance on and usefulness of evidence in both *Standard* and *Summary* modalities suggest that effective summarization does not impact the value of the evidence presented to workers.

6.2 Interaction with Evidence

We now detail the result of the inspection of the number of evidence items participants interacted with in both the *Standard* and *Summary* modalities during the truthfulness assessment tasks. We focus particularly on the number of evidence items clicked by participants for each statement. A specific count of interest was 8, which indicates that participants clicked on exactly one link per

statement (including gold questions), suggesting a uniform and minimal engagement across all statements provided in the task. For the *Summary* modality, we found that out of the total interactions, 511 instances (85.17%) involved clicking exactly eight links to evidence. Similarly, in the *Standard* modality, there were 520 instances (86.67%) where participants engaged with exactly eight links. This close similarity in percentages between the two modalities indicates a comparable level of engagement when workers are provided with either full or summarized content, in the case of workers who do the minimum amount of work. A similar pattern emerges when we exclude instances where participants clicked on exactly eight links. In such cases, the mean number of links clicked in the *Summary* modality was 18.70, with a standard deviation of 8.09, whereas in the *Standard* modality, the mean was slightly lower at 17.50, with a narrower standard deviation of 3.39. This is based on few (approx 15%) of instances, and no statistical significance is measured.

We also extend the analysis to examine the maximum depth of links explored by participants, which reflects how deeply they looked into the provided evidence. We measured the maximum reached depth as the ordinal position of the farthest link clicked (e.g., clicking on the fifth link indicates a max depth of 5). It is again the case that the two modalities exhibit a similar behavior. In the *Summary* modality, 85% of participants clicked no further than the first link, similar to 86.67% in the *Standard* modality. When excluding these cases, participants in the *Summary* modality explored slightly more links on average (mean depth 3.63 vs. 3.43, no statistical significance detected), showing comparable engagement across both modalities.

Table 5: Correlation values between dimensions.

Modality	Need & Usefulness	Need & Time	Usefulness & Time
<i>Standard</i>	0.483	-0.020	0.030
<i>Summary</i>	0.481	0.004	-0.097

6.3 Interaction Across Evaluation Dimensions

In addition to examining the depth of evidence exploration, we also analyzed the number of times each participant engaged with the assessment dimensions (i.e., truthfulness, usefulness of evidence, and need for evidence). Given that there are three dimensions in our evaluation setup, the minimum number of clicks required is three, one for each dimension. This measure helps us understand the decisiveness of participants and the extent to which they reconsider their assessments during the task. Our findings show that in the *Summary* modality, 433 (72.17%) of participants clicked each dimension only once, similar to 450 (75%) in the *Standard* modality. Excluding cases with exactly three clicks, the *Summary* modality shows a mean of 4.45 clicks ($\sigma^2 = 0.67$), slightly lower than the *Standard* modality’s 4.60 clicks ($\sigma^2 = 0.99$). No statistical significance is detected. These statistics suggest that participants in both modalities generally stick to their initial assessments, and the similar levels of engagement suggest that summarized content appears to be sufficient for making assessments. The slightly lower variance in the number of clicks in the *Summary* modality further reinforces this point, suggesting that participants appear slightly more consistently confident in their assessments.

To further explore the interplay between these dimensions, we conduct a correlation analysis examining the relationships among *need*, *usefulness*, and *time elapsed* for both modalities (Table 5). The results reveal a consistent moderate positive correlation ($\rho \approx 0.48$) between *need* and *usefulness*, indicating that participants who perceive the evidence as more necessary also tend to consider it more useful, irrespective of whether they are presented with full-length or summarized content. This relationship underscores the complementary nature of these two dimensions, as both contribute to the assessment process. However, the moderate strength of the correlation suggests that while *need* and *usefulness* are related, they are not interchangeable. Instead, these dimensions capture distinct aspects of participants’ interactions with evidence. On the other hand, the correlations between *need* and *time elapsed*, and *usefulness* and *time elapsed*, are negligible ($\rho \approx -0.02$ to -0.09), showing limited influence of these perceptions on task duration.

To more deeply analyze the interactions among the three assessment dimensions measured for the two modalities (i.e., *time elapsed*, *need*, and *usefulness*), we performed an OLS regression analysis, augmented with ω^2 effect sizes [39], a well-established method in the literature for quantifying effect sizes [18–20, 44, 60]. The results confirm that the *Summary* modality significantly reduces task time when compared to the *Standard* modality ($p < 0.05$, $\omega^2 = 0.0036$). Conversely, perceptions of *need*, *usefulness*, and their interaction exhibited no statistically significant effects on task duration, with negligible ω^2 values (-0.0008 for *need*, -0.0001 for *usefulness*, and -0.0005 for their interaction).

7 Conclusions and Future Work

This study investigates the potential of summarized evidence to enhance the efficiency, effectiveness, and cost-effectiveness of crowdsourced truthfulness assessments. For **RQ1 (Effectiveness)**, our results suggest that the *Summary* modality achieves accuracy and error metrics comparable to the *Standard* modality. While slight differences are observed in MSE and MAE, these are minimal and do not affect the practical utility of the *Summary* approach. These findings affirm that summarization effectively preserves the critical information necessary for accurate truthfulness assessments.


Regarding **RQ2 (Efficiency)**, we find that the *Summary* modality substantially reduces the time required for participants to complete their assessments compared to the *Standard* one. Workers using summarized evidence complete nearly double the number of judgments within the same time frame. Addressing cost-effectiveness, we find that summarization significantly reduces the costs associated with large-scale truthfulness assessments. By enabling workers to complete more judgments within the same time frame, the *Summary* modality reduces overall task costs without reducing quality.

In terms of **RQ3 (Behavioral Insights)**, we observe that the *Summary* modality does not diminish participants’ reliance on or perceived usefulness of evidence. Furthermore, the *Summary* modality produces higher internal agreement among workers, suggesting that the condensed presentation of information aligns participant judgments more effectively. This indicates that summarization not only maintains but also enhances certain aspects of the decision-making process, supporting consistent evidence-based assessments.

Despite these contributions, certain limitations must be acknowledged. First, the reliance on automatically generated summaries introduces the possibility of omitting critical nuances present in full-length evidence, which may impact worker judgment in complex cases. Additionally, our study focuses on a specific dataset and task context, which may limit the generalizability of our findings. Finally, while the *Summary* modality demonstrates promising efficiency, further research is needed to explore its long-term effects on worker fatigue, engagement, and the potential biases.

These findings also suggest directions for future work. We plan to refine summarization techniques to better tailor evidence to specific truthfulness categories or individual worker preferences; also, we plan to test the impact of evidence presentation on worker satisfaction and engagement. Additionally, we plan to generalize these methods to other domains, such as health misinformation or scientific claims. Finally, we plan to explore multi-document summarization techniques that would allow to present multiple pieces of evidence into a single, coherent summary.

Acknowledgments

 This research is partially funded by the European Union (ERC, ExplainYourself, 101077481), as well as a Danish Data Science Academy postdoctoral fellowship (grant number: 2023-1425) and is supported by the Pioneer Centre for AI, DNRF grant number P1. It is also partially supported by the PRIN 2022 Project, “MoT—The Measure of Truth: An Evaluation-Centered Machine-Human Hybrid Framework for Assessing Information Truthfulness”, Code No. 20227F2ZN3, CUP No. G53D23002800006, funded by the European Union – Next Generation EU – PNRR M4 C2 I1.1.

References

- [1] Alim Al Ayub Ahmed, Ayman Aljabouh, Praveen Kumar Donepudi, and Myung Suh Choi. 2021. Detecting Fake News Using Machine Learning: A Systematic Literature Review. *Psychology and Education Journal* 58, 1 (2021), 10 pages. <https://doi.org/10.17762/pae.v58i1.1046>
- [2] Jennifer Allen, Cameron Martel, and David G Rand. 2022. Birds of a Feather Don't Fact-Check Each Other: Partisanship and the Evaluation of News in Twitter's Birdwatch Crowdsourced Fact-Checking Program. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 245, 19 pages. <https://doi.org/10.1145/3491102.3502040>
- [3] 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*, Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault (Eds.). Association for Computational Linguistics, Online, 7352–7364. <https://doi.org/10.18653/v1/2020.acl-main.656>
- [4] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2022. Fact Checking with Insufficient Evidence. *Transactions of the Association for Computational Linguistics* 10 (2022), 746–763. https://doi.org/10.1162/tacl_a_00486
- [5] Isabelle Augenstein, Timothy Baldwin, Meeyoung Cha, Tanmoy Chakraborty, Giovanni Luca Ciampaglia, David Corney, Renee DiResta, Emilio Ferrara, Scott Hale, Alon Halevy, Eduard Hovy, Heng Ji, Filippo Menczer, Ruben Miguez, Preslav Nakov, Dietram Scheufele, Shivam Sharma, and Giovanni Zagni. 2024. Factuality challenges in the era of large language models and opportunities for fact-checking. *Nature Machine Intelligence* 6, 8 (Aug 2024), 852–863. <https://doi.org/10.1038/s42256-024-00881-z>
- [6] Isabelle Augenstein, Christina Lioma, Dongsheng Wang, Lucas Chaves Lima, Casper Hansen, Christian Hansen, and Jakob Grue Simonsen. 2019. MultiFC: A Real-World Multi-Domain Dataset for Evidence-Based Fact Checking of Claims. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 4685–4697. <https://doi.org/10.18653/v1/D19-1475>
- [7] Nadav Borenstein, Greta Warren, Desmond Elliott, and Isabelle Augenstein. 2025. Can Community Notes Replace Professional Fact-Checkers? arXiv. <https://doi.org/10.48550/arXiv.2502.14132> arXiv:2502.14132 [cs.CL]
- [8] Erik Brand, Kevin Roitero, Michael Soprano, Afshin Rahimi, and Gianluca Demartini. 2022. A Neural Model to Jointly Predict and Explain Truthfulness of Statements. *Journal of Data and Information Quality* 15, 1, Article 4 (12 2022), 19 pages. <https://doi.org/10.1145/3546917>
- [9] Alessandro Checco, Kevin Roitero, Eddy Maddalena, Stefano Mizzaro, and Gianluca Demartini. 2017. Let's Agree to Disagree: Fixing Agreement Measures for Crowdsourcing. *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 5, 1 (Sep. 2017), 11–20. <https://doi.org/10.1609/hcomp.v5i1.13306>
- [10] Botambu Collins, Dinh Tuyen Hoang, Ngoc Thanh Nguyen, and Dosam Hwang. 2021. Trends in Combating Fake News On Social Media – A Survey. *Journal of Information and Telecommunication* 5, 2 (2021), 247–266. <https://doi.org/10.1080/24751839.2020.1847379>
- [11] Alphaeus Dmonte, Roland Oruche, Marcos Zampieri, Prasad Callyam, and Isabelle Augenstein. 2025. Claim Verification in the Age of Large Language Models: A Survey. arXiv. <https://doi.org/10.48550/arXiv.2408.14317> arXiv:2408.14317 [cs.CL]
- [12] Xishuang Dong, Shouvon Sarker, and Lijun Qian. 2022. Integrating Human-in-the-loop into Swarm Learning for Decentralized Fake News Detection. In *International Conference on Intelligent Data Science Technologies and Applications (IDSTA)*. IEEE, San Antonio, TX, USA, 46–53. <https://doi.org/10.1109/IDSTA55301.2022.9923043>
- [13] Tim Draws, David La Barbera, Michael Soprano, Kevin Roitero, Davide Ceolin, Alessandro Checco, and Stefano Mizzaro. 2022. The Effects of Crowd Worker Biases in Fact-Checking Tasks. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*. Association for Computing Machinery, Seoul, Republic of Korea, 2114–2124. <https://doi.org/10.1145/3531146.3534629>
- [14] Tim Draws, Alisa Rieger, Oana Inel, Ujwal Gadiraju, and Nava Tintarev. 2021. A Checklist to Combat Cognitive Biases in Crowdsourcing. In *Proceedings of the Ninth AAAI Conference on Human Computation and Crowdsourcing*. AAAI Press, Palo Alto, California, USA, 48–59. <https://doi.org/10.1609/hcomp.v9i1.18939>
- [15] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The Llama 3 Herd of Models. arXiv. <https://doi.org/10.48550/arXiv.2407.21783> arXiv:2407.21783
- [16] Bradley Efron and Robert J. Tibshirani. 1993. *An Introduction to the Bootstrap*. Monographs on Statistics and Applied Probability, Vol. 57. Chapman & Hall/CRC, New York. <https://doi.org/10.1201/9780429246593>
- [17] Carsten Eickhoff. 2018. Cognitive Biases in Crowdsourcing. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (Marina Del Rey, CA, USA) (WSDM '18)*. Association for Computing Machinery, New York, NY, USA, 162–170. <https://doi.org/10.1145/3159652.3159654>
- [18] Nicola Ferro, Yubin Kim, and Mark Sanderson. 2019. Using Collection Shards to Study Retrieval Performance Effect Sizes. *ACM Transactions on Information Systems* 37, 3, Article 30 (March 2019), 40 pages. <https://doi.org/10.1145/3310364>
- [19] Nicola Ferro and Gianmaria Silvello. 2016. A General Linear Mixed Models Approach to Study System Component Effects. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval (Pisa, Italy) (SIGIR '16)*. Association for Computing Machinery, New York, NY, USA, 25–34. <https://doi.org/10.1145/2911451.2911530>
- [20] Nicola Ferro and Gianmaria Silvello. 2018. Toward an anatomy of IR system component performances. *Journal of the Association for Information Science and Technology* 69, 2 (2018), 187–200. <https://doi.org/10.1002/asi.23910>
- [21] David A. Freedman. 2009. *Statistical Models: Theory and Practice* (2 ed.). Cambridge University Press, New York, NY, USA. <https://doi.org/10.1017/CBO9780511815867>
- [22] Meric Altug Gemalmaz and Ming Yin. 2021. Accounting for Confirmation Bias in Crowdsourced Label Aggregation. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI)*. International Joint Conferences on Artificial Intelligence Organization, California, USA, 1729–1735. <https://doi.org/10.24963/ijcai.2021/238>
- [23] Lei Han, Kevin Roitero, Ujwal Gadiraju, Cristina Sarasua, Alessandro Checco, Eddy Maddalena, and Gianluca Demartini. 2019. All Those Wasted Hours: Task Abandonment in Crowdsourcing. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining (Melbourne VIC, Australia) (WSDM '19)*. Association for Computing Machinery, New York, NY, USA, 321–329. <https://doi.org/10.1145/3289600.3291035>
- [24] Linnei Hu, Siqi Wei, Ziwang Zhao, and Bin Wu. 2022. Deep Learning For Fake News Detection: A Comprehensive Survey. *AI Open* 3 (2022), 133–155. <https://doi.org/10.1016/j.aiopen.2022.09.001>
- [25] Christoph Hube, Besnik Fetahu, and Ujwal Gadiraju. 2019. Understanding and Mitigating Worker Biases in the Crowdsourcing Collection of Subjective Judgments. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/3290605.3300637>
- [26] Klaus Krippendorff. 2011. Computing Krippendorff's Alpha-Reliability. https://repository.upenn.edu/asc_papers/43. Technical Report, University of Pennsylvania.
- [27] David La Barbera, Eddy Maddalena, Michael Soprano, Kevin Roitero, Gianluca Demartini, Davide Ceolin, Damiano Spina, and Stefano Mizzaro. 2024. Crowdsourced Fact-checking: Does It Actually Work? *Information Processing & Management* 61, 5 (2024), 103792. <https://doi.org/10.1016/j.ipm.2024.103792>
- [28] David La Barbera, Kevin Roitero, Gianluca Demartini, Stefano Mizzaro, and Damiano Spina. 2020. Crowdsourcing Truthfulness: The Impact of Judgment Scale and Assessor Bias. In *Advances in Information Retrieval*, Joemon M. Jose, Emine Yilmaz, João Magalhães, Pablo Castells, Nicola Ferro, Mário J. Silva, and Flávio Martins (Eds.). Springer International Publishing, Cham, 207–214. https://doi.org/10.1007/978-3-030-45442-5_26
- [29] David M. J. Lazer, Matthew A. Baum, Yoichi Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science* 359, 6380 (2018), 1094–1096. <https://doi.org/10.1126/science.aao2998>
- [30] Huiru Li, Liangxiao Jiang, and Siqing Xue. 2023. Neighborhood Weighted Voting-Based Noise Correction for Crowdsourcing. *ACM Transactions on Knowledge Discovery from Data* 17, 7, Article 96 (4 2023), 18 pages. <https://doi.org/10.1145/3586998>
- [31] Houjiang Liu, Anubrata Das, Alexander Boltz, Didi Zhou, Daisy Pinaroc, Matthew Lease, and Min Kyung Lee. 2024. Human-centered NLP Fact-checking: Co-Designing with Fact-checkers using Matchmaking for AI. *Proceedings of the ACM on Human-Computer Interaction* 8, CSCW2, Article 423 (Nov. 2024), 44 pages. <https://doi.org/10.1145/3686962>
- [32] Houjiang Liu, Jacek Gwizdka, and Matthew Lease. 2024. Exploring Multidimensional Checkworthiness: Designing AI-assisted Claim Prioritization for Human Fact-checkers. arXiv. <https://doi.org/10.48550/arXiv.2412.08185>
- [33] Henry B. Mann and Donald R. Whitney. 1947. On a Test of Whether One of Two Random Variables is Stochastically Larger than the Other. *Annals of Mathematical Statistics* 18, 1 (1947), 50–60. <https://doi.org/10.1214/aoms/1177730491>
- [34] Syed Ishaq Manzoor, Jimmy Singla, and Nikita. 2019. Fake News Detection Using Machine Learning approaches: A systematic Review. In *Proceedings of the 3rd International Conference on Trends in Electronics and Informatics (ICOEI)*. IEEE, Tirunelveli, India, 230–234. <https://doi.org/10.1109/ICOEI2019.8862770>
- [35] Cameron Martel, Jennifer Allen, Gordon Pennycook, and David G. Rand. 2024. Crowds Can Effectively Identify Misinformation at Scale. *Perspectives on Psychological Science* 19, 2 (2024), 477–488. <https://doi.org/10.1177/17456916231190388>
- [36] Preslav Nakov, Alberto Barrón-Cedeño, Giovanni da San Martino, Firoj Alam, Julia Maria Struß, Thomas Mandl, Rubén Miguez, Tommaso Caselli, Mucahid Kutlu, Wajdi Zaghouni, Chengkai Li, Shaden Shaar, Gautam Kishore Shahi, Hamdy Mubarak, Alex Nikolov, Nikolay Babulov, Yavuz Selim Kartal, Michael

- Wiegand, Melanie Siegel, and Juliane Köhler. 2022. Overview of the CLEF-2022 CheckThat! Lab on Fighting the COVID-19 Infodemic and Fake News Detection. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, Alberto Barrón-Cedeño, Giovanni Da San Martino, Mirko Degli Esposti, Fabrizio Sebastiani, Craig Macdonald, Gabriella Pasi, Allan Hanbury, Martin Potthast, Guglielmo Faggioli, and Nicola Ferro (Eds.). Springer International Publishing, Cham, 495–520. https://doi.org/10.1007/978-3-031-13643-6_29
- [37] Preslav Nakov, Giovanni Da San Martino, Tamer Elsayed, Alberto Barrón-Cedeño, Rubén Míguez, Shaden Shaar, Firoj Alam, Fatima Haouari, Maram Hasanain, Nikolay Babulov, Alex Nikolov, Gautam Kishore Shahi, Julia Maria Struß, and Thomas Mandl. 2021. The CLEF-2021 CheckThat! Lab on Detecting Check-Worthy Claims, Previously Fact-Checked Claims, and Fake News. In *Advances in Information Retrieval*, Djoerd Hiemstra, Marie-Francine Moens, Josiane Mothe, Raffaele Perego, Martin Potthast, and Fabrizio Sebastiani (Eds.). Springer International Publishing, Cham, 639–649. https://doi.org/10.1007/978-3-030-72240-1_75
- [38] Thanh Tam Nguyen, Matthias Weidlich, Hongzhi Yin, Bolong Zheng, Quang Huy Nguyen, and Quoc Viet Hung Nguyen. 2020. FactCatch: Incremental Pay-as-You-Go Fact Checking with Minimal User Effort. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Virtual Event, China) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 2165–2168. <https://doi.org/10.1145/3397271.3401408>
- [39] Stephen F. Olejnik and James Algina. 2003. Generalized Eta and Omega Squared Statistics: Measures of Effect Size for Some Common Research Designs. *Psychological Methods* 8, 4 (2003), 434–447. <https://doi.org/10.1037/1082-989X.8.4.434>
- [40] Gordon Pennycook and David G. Rand. 2019. Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences* 116, 7 (2019), 2521–2526. <https://doi.org/10.1073/pnas.1806781116>
- [41] Gordon Pennycook and David G. Rand. 2021. The Psychology of Fake News. *Trends in Cognitive Sciences* 25, 5 (2021), 388–402. <https://doi.org/10.1016/j.tics.2021.02.007>
- [42] Yunke Qu, Kevin Roitero, David La Barbera, Damiano Spina, Stefano Mizzaro, and Gianluca Demartini. 2022. Combining Human and Machine Confidence in Truthfulness Assessment. *Journal of Data and Information Quality* 15, 1, Article 5 (dec 2022), 17 pages. <https://doi.org/10.1145/3546916>
- [43] Benjamin Riedel, Isabelle Augenstein, Georgios P Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the fake news challenge stance detection task. arXiv. <https://doi.org/10.48550/arXiv.1707.03264>
- [44] Kevin Roitero, Ben Carterette, Rishabh Mehrotra, and Mounia Lalmas. 2020. Leveraging Behavioral Heterogeneity Across Markets for Cross-Market Training of Recommender Systems. In *Companion Proceedings of the Web Conference 2020* (Taipei, Taiwan) (WWW '20). Association for Computing Machinery, New York, NY, USA, 694–702. <https://doi.org/10.1145/3366424.3384362>
- [45] Kevin Roitero, Michael Soprano, Shaoyang Fan, Damiano Spina, Stefano Mizzaro, and Gianluca Demartini. 2020. Can The Crowd Identify Misinformation Objectively? The Effects of Judgment Scale and Assessor's Background. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Xi'an, China (Virtual)) (SIGIR '20). Association for Computing Machinery, New York, NY, USA, 439–448. <https://doi.org/10.1145/3397271.3401112>
- [46] Mohammed Saeed, Nicolas Traub, Maelle Nicolas, Gianluca Demartini, and Paolo Papotti. 2022. Crowdsourced Fact-Checking at Twitter: How Does the Crowd Compare With Experts?. In *Proceedings of the 31st ACM International Conference on Information and Knowledge Management* (Atlanta, GA, USA) (CIKM '22). Association for Computing Machinery, New York, NY, USA, 1736–1746. <https://doi.org/10.1145/3511808.3557279>
- [47] Michael Schlichtkrull, Yulong Chen, Chenxi Whitehouse, Zhenyun Deng, Mubashara Akhtar, Rami Aly, Zhijiang Guo, Christos Christodoulopoulos, Oana Cocarascu, Arpit Mittal, James Thorne, and Andreas Vlachos (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 1–26. <https://doi.org/10.18653/v1/2024. fever-1.1>
- [48] Vinay Setty. 2024. Surprising Efficacy of Fine-Tuned Transformers for Fact-Checking over Larger Language Models. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington DC, USA) (SIGIR '24). Association for Computing Machinery, New York, NY, USA, 2842–2846. <https://doi.org/10.1145/3626772.3661361>
- [49] Michael Soprano, Kevin Roitero, David La Barbera, Davide Ceolin, Damiano Spina, Gianluca Demartini, and Stefano Mizzaro. 2024. Cognitive Biases in Fact-Checking and Their Countermeasures: A Review. *Information Processing & Management* 61, 3 (2024), 103672. <https://doi.org/10.1016/j.ipm.2024.103672>
- [50] Michael Soprano, Kevin Roitero, David La Barbera, Davide Ceolin, Damiano Spina, Stefano Mizzaro, and Gianluca Demartini. 2021. The Many Dimensions of Truthfulness: Crowdsourcing Misinformation Assessments on a Multi-dimensional Scale. *Information Processing & Management* 58, 6 (2021), 102710. <https://doi.org/10.1016/j.ipm.2021.102710>
- [51] All Abdullah Tanvir, Ehasas Mia Mahir, Saima Akhter, and Mohammad Rezwani Huq. 2019. Detecting Fake News using Machine Learning and Deep Learning Algorithms. In *Proceedings of the 7th International Conference on Smart Computing & Communications* (ICSCC). IEEE, Piscataway, NJ, USA, 1–5. <https://doi.org/10.1109/ICSCC.2019.8843612>
- [52] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a Large-scale Dataset for Fact Extraction and VERification. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Marilyn Walker, Heng Ji, and Amanda Stent (Eds.). Association for Computational Linguistics, New Orleans, Louisiana, 809–819. <https://doi.org/10.18653/v1/N18-1074>
- [53] James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal. 2019. The FEVER2.0 Shared Task. In *Proceedings of the Second Workshop on Fact Extraction and Verification (FEVER)*, James Thorne, Andreas Vlachos, Oana Cocarascu, Christos Christodoulopoulos, and Arpit Mittal (Eds.). Association for Computational Linguistics, Hong Kong, China, 1–6. <https://doi.org/10.18653/v1/D19-6601>
- [54] Nikhita Vedula and Srinivasan Parthasarathy. 2021. FACE-KEG: Fact Checking Explained using Knowledge Graphs. In *Proceedings of the Second International Conference on Web Search and Data Mining* (Virtual Event, Israel) (WSDM '21). Association for Computing Machinery, New York, NY, USA, 526–534. <https://doi.org/10.1145/3437963.3441828>
- [55] Nguyen Vo and Kyumin Lee. 2019. Learning from Fact-checkers: Analysis and Generation of Fact-checking Language. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR '19). Association for Computing Machinery, New York, NY, USA, 335–344. <https://doi.org/10.1145/3331184.3331248>
- [56] Yuxia Wang, Revanth Gangi Reddy, Zain Muhammad Mujahid, Arnab Arora, Aleksandr Rubashevskii, Jiahui Geng, Osama Mohammed Afzal, Liangming Pan, Nadav Borenstein, Aditya Pillai, Isabelle Augenstein, Iryna Gurevych, and Preslav Nakov. 2024. Factcheck-Bench: Fine-Grained Evaluation Benchmark for Automatic Fact-checkers. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 14199–14230. <https://doi.org/10.18653/v1/2024.findings-emnlp.830>
- [57] Greta Warren, Irina Shklovski, and Isabelle Augenstein. 2025. Show Me the Work: Fact-Checkers' Requirements for Explainable Automated Fact-Checking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (CHI '25). ACM, Honolulu, HI, USA, 1–21. <https://doi.org/10.1145/3706598.3713277>
- [58] Jing Yang, Didier Vega-Oliveros, Tais Seibt, and Anderson Rocha. 2021. Scalable Fact-checking with Human-in-the-Loop. In *2021 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, Montpellier, France, 1–6. <https://doi.org/10.1109/WIFS53200.2021.9648388>
- [59] Barry Menglong Yao, Aditya Shah, Lichao Sun, Jin-Hee Cho, and Lifu Huang. 2023. End-to-End Multimodal Fact-Checking and Explanation Generation: A Challenging Dataset and Models. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Taipei, Taiwan) (SIGIR '23). Association for Computing Machinery, New York, NY, USA, 2733–2743. <https://doi.org/10.1145/3539618.3591879>
- [60] Fabio Zampieri, Kevin Roitero, J. Shane Culpepper, Oren Kurland, and Stefano Mizzaro. 2019. On Topic Difficulty in IR Evaluation: The Effect of Systems, Corpora, and System Components. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval* (Paris, France) (SIGIR '19). Association for Computing Machinery, New York, NY, USA, 909–912. <https://doi.org/10.1145/3331184.3331279>
- [61] Andy Zhao and Mor Naaman. 2023. Variety, Velocity, Veracity, and Viability: Evaluating the Contributions of Crowdsourced and Professional Fact-checking. <https://doi.org/10.31235/osf.io/yfxd3>