

Evaluation of Reliability Criteria for News Publishers with Large Language Models

Manuel Pratelli*

IIT-CNR

Pisa, Italy

IMT School for Advanced Studies Lucca

Lucca, Italy

manuel.pratelli@imtlucca.it

Fabio Pinelli

IMT School for Advanced Studies Lucca

Lucca, Italy

fabio.pinelli@imtlucca.it

John Bianchi

IMT School for Advanced Studies Lucca

Lucca, Italy

john.bianchi@imtlucca.it

Marinella Petrocchi

IIT-CNR

Pisa, Italy

IMT School for Advanced Studies Lucca

Lucca, Italy

marinella.petrocchi@iit.cnr.it

Abstract

In this study, we investigate the use of a large language model to assist in the evaluation of the reliability of the vast number of existing online news publishers, addressing the impracticality of relying solely on human expert annotators for this task. In the context of the Italian news media market, we first task the model with evaluating expert-designed reliability criteria using a representative sample of news articles. We then compare the model's answers with those of human experts. The dataset consists of 352 news articles annotated by three human experts and the LLM. Examining 6,081 annotations over six criteria, we observe good agreement between LLM and human annotators in three evaluated criteria, including the critical ability to detect instances where a text negatively targets an entity or individual. For two additional criteria, such as the detection of sensational language and the recognition of bias in news content, LLMs generate fair annotations, albeit with certain trade-offs. Furthermore, we show that the LLM is able to help resolve disagreements among human experts, especially in tasks such as identifying cases of negative targeting.

CCS Concepts

• **Information systems** → **Information retrieval**; • **Human-centered computing** → **Social tagging systems**; • **Computing methodologies** → **Information extraction**.

Keywords

Reliability Evaluation, Good Editorial Practices, Generative Question Answering, LLMs, Inter-annotator agreement

*Corresponding Author



This work is licensed under a Creative Commons Attribution 4.0 International License. *Websci '25, New Brunswick, NJ, USA*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1483-2/25/05

<https://doi.org/10.1145/3717867.3717924>

ACM Reference Format:

Manuel Pratelli, John Bianchi, Fabio Pinelli, and Marinella Petrocchi. 2025. Evaluation of Reliability Criteria for News Publishers with Large Language Models. In *Proceedings of the 17th ACM Web Science Conference 2025 (Websci '25)*, May 20–24, 2025, New Brunswick, NJ, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3717867.3717924>

1 Introduction

The decline of the traditional journalism system, combined with difficulties in maintaining editorial control, has raised concerns about the quality of information disseminated by online news publishers. Specialized organizations [2, 5, 7–9], have been instrumental in assessing the reliability of many online news sources.




Figure 1: Traditional Approach to Evaluating the Reliability of a News Publisher

Each organisation conducts the evaluation process by defining a specific set of criteria and identifying the expert profiles required for their assessment (e.g., journalists or trained annotators). Criteria typically consider elements like the inclination to publish propaganda or politically biased content [18, 26]. The phase of the criteria evaluation process generally lacks automation. This process can be modeled as depicted in Figure 1. The organization in charge of evaluating publishers sets criteria for "good journalism" and assigns experts in the field to manually evaluate the publisher against these criteria. The evaluation results are aggregated at the criterion level, taking into account potential disagreements among experts, to produce a final score - often numerical in nature. This score serves as a holistic measure of the publisher's overall reliability, providing a comprehensive assessment without detailing

the specific contributions of individual criteria. Although these scores offer valuable insights [21, 21, 26, 38], the process of evaluating individual news outlets is labor-intensive and time-consuming. Manual assessments, where expert annotators scrutinize ownership information and content, remain crucial yet demanding [23, 34].

In recent years, Large Language Models (LLMs) have excelled at generating text that closely mimics human language [25]. Their capabilities extend to various natural language processing tasks, such as sentiment analysis and text summarization [40]. This capability has paved the way for numerous potential applications, including improving educational processes and providing answers to medical questions [17].

On the one hand, we have an accurate but time-consuming evaluation process: it takes months to evaluate a single news source, starting with the selection and training of annotators, through the actual evaluation of specific criteria, and ending with the calculation of the final score, not to mention that after months the situation of the media outlet itself may have changed and require a new analysis. On the other hand, a generative intelligence whose evolution seems to progress constantly, see, e.g., the recent release of GPT-4o (omni)[10].




Figure 2: The core idea: Given a set of criteria and a set of news articles, we evaluate the agreement between the manual evaluation of the criteria and the automated evaluation to determine whether automated evaluation is a valid approach to assessing reliability of the articles’ publishers.

Objective and Approach. The main goal of this work is to develop a more scalable and more efficient trust evaluation process by incorporating automation into the traditional approach of manually evaluating good journalism criteria defined by specialized organizations, see Figure 1. To the best of our knowledge, this is the first study that proposes to introduce automation into the task of evaluating such criteria. To conduct the experiment, as shown in Figure 2 that represent the core idea of the work, we use a two-step process. In the first step, we perform the two following tasks in parallel: (i) the selection and implementation of a set of quality criteria, and (ii) the collection of a representative sample of publishers and news articles. In the second step, three experienced human annotators evaluate the criteria on each of the sampled articles. Their answers are compared with those of an LLM. The intuition underlying this study is that agreement between human experts and the LLM on individual criteria leaves room for automation of related subtasks. In addition, we explore the potential of the LLM to assist in resolving disagreements among human experts.

Research Questions. Achieving our goal translates into answering the following questions:

- **RQ1:** How to define quality prompts to help LLMs correctly evaluate good journalism criteria to assess the reliability level of a publisher?
- **RQ2:** How good is an LLM at aligning with the responses given by human-experts when evaluating the good journalism criteria?
- **RQ3:** What are the types of questions for which an LLM can effectively support, dare we say even replace, an expert annotator in the process of evaluating the reliability of a publisher?

Results. Through a process of refinement, we are able to write prompts that are optimized so that 1) the LLM understands the questions it needs to answer, and 2) the answers are consistent with those of human annotators. The results show agreement between the LLM and the annotators in recognizing the type of the news, the presence of a fact-based *lede* at the beginning of the article and whether the text negatively targets someone or something. Also, the evaluation of labor-intensive criteria, such as checking for the presence of article bias and sensational language, achieves a fair level of agreement, albeit with certain trade-offs. In addition, the LLM can provide valuable support in cases of disagreement between human experts in different evaluation tasks.

Applications. The ability to automatically generate high-quality, criteria-specific responses from the analysis of news articles opens up at least two potential avenues of application:

- On an organizational level, automated evaluation of criteria can be integrated into existing workflows for assessing the reliability of news publishers. This is particularly effective in resource-constrained environments where automation provides scalability. For example, assessing labor-intensive criteria on a large volume of news articles is often not feasible due to limited expertise or funding. Automation, on the other hand, enables large-scale, near-real-time monitoring and analysis, significantly improving the efficiency and reach of reliability assessments.
- At the user level, criteria-specific responses can increase readers’ awareness of the content they consume in real time. The goal is to highlight potential deviations in specific aspects (e.g., sensational language) from established quality standards of good journalism. This approach translates expert practices into tools accessible to end-users, fostering critical thinking and enabling more informed and reflective engagement with news content.

2 Related Work

Towards Automating the Evaluation of News Source Reliability: Challenges, Advances, and Gaps. Several studies have proposed lists of criteria to distinguish reliable news sites from unreliable ones. Notable examples include criteria developed by: (i) independent organizations, such as the Global Disinformation Index (GDI) [5], NewsGuard [9], and Media Bias Fact Check (MBFC) [8]; (ii) platforms, such as Facebook’s “Tips to Spot False News” [4]; (iii) academic research [24, 41]; and (iv) collaborative projects [3, 13].

The criteria are designed for different applications and annotator profiles, ranging from lay users to expert human annotators. As a result, the granularity of the evaluation and the complexity of

the tasks required vary. For example, Facebook’s guidelines [4] provide simple tips to help ordinary users identify misinformation and make informed decisions. In contrast, the criteria developed by GDI [5] and NewsGuard [31] are tailored for expert annotators - such as trained professionals or journalists. While these criteria are clear and interpretable, applying them to a news article often requires expertise in the field. In other words, experts and non-experts (e.g., crowd raters) might disagree [19]. For example, verifying the presence of bias in the textual content of news articles might not be trivial for a common user to evaluate.

In recent years, substantial evidence has emerged supporting the effectiveness of professional and user-generated ratings applied at the text level [21] or source level [21, 26, 38] in mitigating the spread of low-quality content. In particular, ratings based on expert criteria show stronger effects in reducing the intention to share low-quality news [21]. However, manually evaluating content at the professional level, such as fact-checking individual news headlines [21], poses significant scalability challenges. Even more, continuous evaluation of source-level ratings, which often aggregate ratings of article-level criteria [5, 31], is highly impractical given the massive volume and rapid distribution of news content online.

Given the challenges of scalability and the recognition of source rating estimation as a key intervention to mitigate the spread of low-quality content [27], we highlight the paucity of research addressing these critical needs.

Some recent work aims to automate the evaluation of the reliability of news sources, exploring different approaches. One method relies solely on user interactions within social media discussions, ignoring the content of the articles themselves [37]. Another approach infers the reliability of the news publisher by classifying the content of individual articles [20]. A third line of research involves using LLMs to directly evaluate news outlets [39].

With specialized organizations in mind, the Global Disinformation Index (GDI) uses LLMs to encode tagged phrases from known disinformation content, creating digital filters. These filters are applied to new material to flag sites for manual review if a significant number of sentences match established disinformation patterns. A team of trained analysts conducts weekly manual reviews [32]. To our knowledge, GDI does not automate the process of criteria evaluation. Instead, it uses LLM to assess similarity with analyst-generated narrative filters that serve as examples of disinformation themes. NewsGuard uses artificial intelligence in proprietary machine learning processes to identify false claims circulating online in real time [33].

The closest work to ours that we found by studying the literature is that of Yang and Menczer [39]. It considers more than 7,000 news outlets, in a variety of languages, and asks ChatGPT 3.5 (March 1, 2023 release) to rate each of these news outlets from 0 to 1 (0 no credibility, 1 max credibility), giving as input only the name of the news outlet. The result was then compared with the ratings given to them NewsGuard and the Media Bias Fact Check, as well as by aggregate human ratings produced by Lin and others [29]. The ratings given by ChatGPT correlate with those of the expert annotators (Spearman’s $\rho = 0.54$, $p < 0.001$). Although one of the goals of the present work is similar to that of [39], i.e., to evaluate the agreement between the LLM and expert annotators

regarding the reliability of online news outlets, the two studies differ in methodology. In terms of methodology, Yang and Menczer’s prompt is concise and high-level, asking the GPT to rate the reliability of the media outlet within a range of values. Our prompt is more complex: we ask the LLM to rate 6 criteria out of 352 news articles, and we compare all LLM responses with those of human annotators. The capillarity of the study (asking the LLM a series of specific questions) is preparatory to understanding which steps in the process can be automated through the use of LLMs. Thus, to the best of our knowledge, no prior work has proposed a methodology for using LLMs to automate the evaluation of a complete set of widely accepted, expert-designed criteria for assessing the quality of online information.

We identify two key gaps in the current scientific literature: (i) the lack of a generally accepted list of evaluation criteria, and (ii) the lack of a comprehensive ground truth dataset to support research in this area. Focusing on the Italian media landscape, this study represents the first attempt to address these gaps. Specifically, we: (i) analyze a set of "worthy-to-automate" criteria: we select them because they are adopted for conducting real-world news website reliability evaluation from two expert-designed frameworks: NewsGuard [31] and the Global Disinformation Index (GDI) (see Table 1); (ii) explore the potential of LLMs to automate their evaluation; and (iii) provide a resulting annotated dataset to facilitate further research in this area.

3 Methodology

This section presents our methodology. The goal is to define appropriate prompts for the LLM to evaluate journalistic criteria used to assess the reliability of online publishers, and measure the agreement between the LLM’s answers and those of experienced annotators.

3.1 Framework overview

We start with the selection (Section 3.2) and implementation (Section 3.3) of a set of criteria for evaluating the reliability of online news publishers (upper left blue dotted box in Figure 2). As explained in 3.3, in the context of this study, criteria implementation means defining a list of questions that can be evaluated by both a human expert and an LLM-based annotator. Each question aims to highlight a specific aspect of a news item according to the corresponding criterion. Second, as described in Section 3.4, we proceed to select a set of publishers to evaluate and extract a sample of news articles from those publishers (lower left green dotted box in Figure 2). Then, in 3.5, we ask three experienced annotators and the LLM to answer the questions representing the criteria. Finally, in the *Results* section, in case of agreements between human experts we evaluate the agreement between the experts’ answer and the LLM answer (right red box in Figure 2). In cases where there is disagreement between experts, we explore the potential of the LLM to facilitate resolution of these conflicts.

3.2 Criteria Selection

The goal of this section is to select a *worthy-to-automate* list of criteria. As mentioned earlier, these criteria must be both widely recognized and labor intensive to evaluate manually.

The starting point for this analysis is the list of criteria proposed by [36]. This work identifies, from the complete list of GDI criteria [35], a preliminary set of *worthy-to-automate* criteria. A detailed explanation of these criteria can be found in Table 1. These criteria are selected because, although they have different names, and their definitions do not match NewsGuard’s definition, there is a strong conceptual correspondence between them [36]. The NewsGuard list of criteria is particularly noteworthy, as it is widely used [24]) too. This consistency in definition across both organizations assures that these criteria are widely accepted as effective for conducting real-world reliability assessments of news publishers.

The tests in [36] were conducted on the Italian media market, which is also the focus of this study. For the scope of this paper, however, we have excluded some criteria from the original set. In particular, we will not consider in our analysis the criteria that do not involve the analysis of a written text (like *VisPres*) and the (trivial-to-evaluate) features that can be extracted using heuristic methods (*ByInfo*). Also, *Common Coverage* is excluded due to their reliance on multi-source analysis. The last six criteria in Table 1 cover broader editorial and operational aspects of news outlets, such as editorial independence, funding sources, and ownership structure. These require access to meta-information that is not typically part of the article text, making them less amenable to automated evaluation by an LLM.

The resulting set of *worthy-to-automate* criteria is as follows:

- *Headline Accuracy*: Assesses the congruence between a headline and the corresponding article content.
- *Lede Presence*: Checks for the presence of a factual summary at the beginning of an article.
- *Negative Targeting*: Evaluates whether the article negatively targets specific individuals or institutions.
- *Article Bias*: Gauges the fairness and balance of the article’s content.
- *Sensational Language*: Identifies emotionally charged or exaggerated language that could mislead readers.

3.3 Criteria Implementation

In this paper, implementing a criterion means defining a clear and concise question, along with potential responses, that can be effectively evaluated by both human experts and LLM-based annotators, using the text of a news article as input. Table 2 details how we have implemented the criteria selected in Section 3.2. Most questions closely align with the descriptions of the criteria provided by GDI. For example, GDI criterion *Negative Targeting*: “*Rating of whether the story negatively targets a specific individual or group*” is made concrete by the question: “*Does the article negatively target individuals or groups?*”. We also instruct the model to specify the reason why an individual or group is negatively targeted (see “*NegTarg (Identification)*” in Table 2). This refinement allows us to assess the ability of the model to identify specific groups or minorities that are negatively targeted in the news. Throughout the manuscript, the general term “*NegTarg*” refers to “*NegTarg (Detection)*”, which denotes the ability to detect textual content that negatively targets groups or individuals. In Section 5.2, we explicitly distinguish between Detection and Identification responses.

Table 1: GDI criteria with their definitions

Criteria Name	Short name	Definition
Headline accuracy	HeadAcc	Rating for how accurately the story’s headline describes the content of the story. Indicative of clickbait
Byline information	ByInfo	Rating for how much information is provided in the article’s byline. Attribution of stories creates accountability for their veracity
Lede present	LedePres	Rating for whether the article begins with a fact-based lede. Indicative of fact-based reporting and high journalistic standards
Common coverage	ComCov	Rating for whether the same event has been covered by at least one other reliable local media outlet. Indicative of a true and significant event
Recent coverage	RecCov	Rating for whether the story covers a news event or development that occurred within 30 days prior to the article’s publication date. Indicative of a newsworthy event, rather than one which has been taken out of context
Negative targeting	NegTarg	Rating for whether the story negatively targets a specific individual or group. Indicative of hate speech, bias or an adversarial narrative.
Article bias	ArtBias	Rating for the degree of bias in the article. Indicative of neutral fact-based reporting or well-rounded analysis.
Sensational language	SensLang	Rating for the degree of sensationalism in the article. Indicative of neutral fact-based reporting or well-rounded analysis.
Visual presentation	VisPres	Rating for the degree of sensationalism in the visual presentation of the article. Indicative of neutral fact-based reporting or well-rounded analysis.
Attribution	Attr	Rating for the number of policies and practices identified on the site. Assesses policies regarding the attribution of stories, facts, and media (either publicly or anonymously); indicative of policies that ensure accurate facts, authentic media and accountability for stories.
Comment policies	CommPol	Rating for the number of policies identified on the site. Assesses policies to reduce disinformation in user-generated content. Rating for the mechanisms to enforce comment policies identified on the site. Assesses the mechanism to enforce policies to reduce disinformation in user-generated content
Editorial principles and practices	EdPrincPract	Rating for the number of policies identified on the site. Assesses the degree of editorial independence and the policies in place to mitigate conflicts of interest. Rating for the degree to which the site is likely to adhere to an ideological affiliation, based on its published editorial positions. Indicative of politicized or ideological editorial decision-making
Ensuring accuracy	EnsAcc	Rating for the number of policies and practices identified on the site. Assesses policies to ensure that only accurate information is reported.
Funding	Fund	Rating for the number of revenue sources identified on the site. Indicative of possible conflicts of interest stemming from over-reliance on one or few sources of revenue
Ownership	Own	Rating based on the number of distinct executive- or board-level financial and editorial decision-makers listed on the site. Indicative of a separation between financial and editorial decision-making, to avoid conflicts of interest.

In addition to the selected criteria reported in section 3.2, we also introduce a meta-criterion (*Type*) to assess whether the LLM can accurately determine the type of article it is analyzing, e.g., distinguishing between straight news and other journalistic forms such as editorials, satire or soft news [15].

3.4 Articles Dataset

This study focuses on the Italian media market. To ensure an unbiased analysis, it is essential to gather textual data from articles that best represent the Italian online media landscape. The approach adopted involves collecting articles from a list of representative Italian news publishers. Specifically, a past study about the Italian online news market [35] proposed a balanced list of Italian news publishers, considering: (i) the media outlet distribution (national or local), (ii) the geographic location of the media outlet, (iii) the political orientation, and (iv) the disinformation risk scores. Table 5 (Appendix C) shows the online publishers selected in [35]. We have included these 34 publishers in this study because they have already been selected to be representative of a country’s news market.

Furthermore, for the current analysis we have collected news articles that were published at the same time as the articles in the original study. This is because the quality of a newspaper can vary over time, and we wanted to try to replicate the same distribution of news articles, both by source and risk of exposure to disinformation. Thus, for each of the 34 news publishers, we retrieved a mean of 10 articles published within a 7-month period, from April to October 2021 for a total of 352 news articles, which is consistent with the timeframe and sample size used in the GDI analysis.

We would like to emphasize that the primary goal of this study is not to evaluate the reliability of a single publisher. Instead, our focus is on assessing the ability of a large language model to evaluate specific journalistic quality criteria applied to a curated collection of article texts.

Table 2: Selected criteria with their short names and implementations

Criteria Name	Short name	Implementation	Answers' Options
Headline accuracy	HeadAcc	How accurate is the news's headline with the content of the news?	Inaccurate/Quite inaccurate/Quite accurate/Accurate
Lede present	LedePres	Does the article start with a summary of the main facts?	Yes/No
Negative targeting	NegTarg	(<i>Detection</i>) Does the article negatively target individuals or groups? (<i>Identification</i>) Indicate what issue the group or individual is negatively targeted on.	Yes/No issue (Politics / Gender / Religion / Other)
Article bias	ArtBias	How much biased is the article?	Biased/Quite biased/Quite unbiased/Unbiased
Sensational language	SensLang	How sensational is the tone of the news?	Sensational/Quite sensational/Quite neutral/Neutral
Meta-criterion			
Type of news	Type	What kind of news are you reading?	Straight news/Editorial/Investigation/Satire/Soft News

The 352 article texts were collected using an automated pipeline combining several tools. First, we employed the Selenium library [12] to gather URLs from publishers' websites. Next, we utilized the GNU Wget command [14] to download each article's HTML page. Finally, XPATH queries were used to extract the textual content, including the title and main text of each article.

3.5 Criteria Evaluation

This section introduces the human annotators and explains their annotation process. In addition, the section presents the design of the LLM-annotator prompts.

Human Expert Annotations. Initially, we sanitized the text of the articles to ensure anonymity of publishers and authors to prevent biasing the annotators' judgments. Then, we asked two experts from the Media and Communication Unit of our department, and a Ph.D. student in Data Science, to annotate the articles. Articles, questions, and options for answers were uploaded to the Google Drive platform. Each article was annotated by two of the three annotators. The annotators were able to consult with each other when in doubt: This is because we want their annotation to be a ground truth against which we can later compare LLM responses.

Prompt Design. The prompt design process ends with the formulation of the questions shown in Table 2. For each of the articles, we ask the LLM the questions listed in the table (column *Implementation*) and give some options for the answers (column *Answers' Options*). It is important to note that since our dataset consists of Italian-language news, the questions we pose to the LLM and the answers we expect are in Italian. In Table 2, we have provided the English translation for the convenience of the reader. The validity of the answers is compared to the answers of the expert annotators, which we consider to be ground truth. The agreement, measured by Cohen's Kappa [22], between the model's answers and those of the experts provides a measure of the quality of the prompts.

During the prompt design phase, we analyze the cases with the most disagreements and identify three main motivations: (i) the prompt wording is not adequate for correct processing by the LLM, (ii) there are errors in the experts' annotations, and (iii) the LLM has difficulty interpreting certain contexts or questions correctly.

For prompts categorized under (i), we reviewed the initial versions of questions and answers until the latter were consistent with one of the answer options. Consistency was evaluated on a selected subset of news articles.

In modifying the prompts, both at the *question* level and the *answer* level, we preserved the semantics so that the validity of the answers given by human annotators remain unchanged, while improving the model's interpretive capabilities. Table 3 shows how we refined the prompts.

The *HeadAcc* criterion was rephrased to enhance clarity. Similarly, the question for the *LedePres* criterion was revised to explicitly incorporate the definition of an article's *lede* within the question itself.

Some response options remain unchanged. Specifically, *HeadAcc*, *LedePres*, and *NegTarg* continue to be binary responses. For *ArtBias* and *SensLang*, as detailed in Section 5.1, converting the original responses into binary formats enhances the consistency and reliability of the scoring process. Regarding the *Type* meta-criterion, modifying the term "*Editorial*" to "*Opinion*" has proven to improve the quality of the responses.

After testing the new prompts on a subset of the news articles and manually checking the consistency of LLM's responses, we repeated the experiment on the entire dataset. In the cases of *ArtBias* and *SensLang*, to implement the changes outlined in Table 3, we reduced the original four-class responses to two classes. This remapping allows us to maintain both datasets without the need for additional data collection.

During the prompt refinement process, we addressed the cases that fell under category (i). The improvements achieved through this process and the ability of the LLM to align with human-annotators are detailed in section 5.1. For cases categorized under (ii) and (iii), the observed issues stemmed from either human annotation errors or the intrinsic limitations of the model. In Section 5.2, we investigate scenarios where the LLM can support in resolving disagreements between experts.

Collection of LLM annotations: The model is instructed to assume the persona of an experienced journalist to ensure that the answers are appropriate for that role. Each question was asked to LLM three times, this was to check that it was not answering inconsistently,

Table 3: Criteria implementation refinement

Short Name	Question-Level Modification	Response-Level Modification	Description
HeadAcc	Rephrased for clarity	None	-
LedePres	Rephrased for clarity	None	Introduced the definition of lede into the question
NegTarg	None	None	-
ArtBias	None	Response options simplified to Biased/Unbiased	-
SensLang	None	Response options simplified to Sens/Neutral	-
Meta-criterion			
Type	None	Response option changed from "Editorial" to "Opinion"	-

perhaps as a result of hallucinations. The answers given by LLM were always the same for the three rounds, thus, a single answer was produced for each article and criterion.

Annotated dataset of human and LLM responses: At the end of this process, we obtained a total of 6,081 annotations across the six evaluation criteria (see Section 3.2) for the dataset of 352 news articles (see Section 3.4). In addition, as described in Section 3.3, we conducted a more detailed analysis of one specific criterion—negative targeting—by introducing target identification. This refinement resulted in an additional 1,011 annotations (see “NegTarg (Identification)” in Section 5.2).

4 Experimental Setup

Agreement Metric: We use Cohen’s Kappa [22] to measure the agreement between annotators. The formula to calculate Cohen’s Kappa is as $k = (p_o - p_e) \div (1 - p_e)$ where p_o is the observed probability of agreement between the annotators and p_e is the expected probability of agreement if the annotators assigned the labels randomly [16]. We calculate this metric using the implementation available in the scikit-learn library[11]. As for the interpretation of values, we refer to the descriptions by Mary L. McHugh [30], which are based on Jacob Cohen’s original interpretations. A kappa value between 0 and 0.20 indicates no agreement among raters, beyond what would be expected by chance. Values from 0.21 to 0.39 represent minimal agreement. A weak level of agreement is observed for values between 0.40 and 0.59, while values between 0.60 and 0.79 indicate a moderate agreement. A strong agreement is denoted by kappa values from 0.80 to 0.90. Finally, values above 0.90 suggest an almost perfect agreement between the raters.

Model selection and configuration: As LLM, we selected ChatGPT-4o, the most recent model released by OpenAI at the time of the experiments [6]. We configure the model’s hyperparameters to optimize the quality of its output. The temperature is set to zero because we need consistent and predictable responses. The *Maximum Length* parameter has not been used because article lengths vary widely and a fixed limit would not be appropriate in all cases. Both *Frequency Penalty* and *Presence Penalty* are set to their default values (0) because avoiding repetition is not a priority for our task. The *Stop Sequences* parameter is set to *None* because we do not need to stop text generation at a particular end sequence.

5 Results

First, we see how well the LLM’s responses match those of the human annotators, considering the original prompts in Table 2. Then, in cases of strong disagreement, we see if using the refined

prompts (Table 3) makes things better. As a third analysis, we look at the cases where the human annotators disagree to see if the LLM could provide support.

5.1 Agreement between experts and LLM




Figure 3: The average agreement between experts and the LLM (computed considering only cases in which the experts themselves reached a consensus).

In Figure 3, we report the agreement between LLM and human annotators *for only those cases where human annotators agree with each other*. Essentially, for each article and criterion, the single annotation value determined by human experts (equal to the two concordant annotation values) is compared with the response provided by the LLM.

For 3 of the 6 criteria analyzed, the agreement between LLM and experts exceeds (or equals, in the case of LedePres) the minimum acceptance threshold set at 0.4. On the other hand, when assessing the degree of sensationalism in the article or the presence of bias, LLM disagrees with the expert annotators, who let us remember that it is the ground truth for us. Specifically, we obtain $k = 0.2064$ for *ArtBias* and $k = 0.1732$ for *SensLang*. We get the highest value of the agreement for *NegTarg*, $k = 0.7089$. Without prompt modification, the LLM is very careful to detect whether a news text is negatively targeting a group or a person.

We went on to analyze in more detail what happens to the LLM when it has to evaluate the criteria *SensLang* and *ArtBias*.

Analyzing the confusion matrix of *SensLang* (Figure 4) and *ArtBias* (Figure 5), we see that the LLM is more sensitive (i.e., it detects sensational language and bias where human annotators do not);

however, it makes errors mainly in neighboring classes. In both confusion matrices, 1 means maximum degree of sensationalism/bias, and 4 means no sensationalism/no bias. Looking at the case of sensationalism, for example, we see that for experts, the articles read are mostly neutral (10+96+82), while for LLM, only (82+2) have no trace of sensationalism. This is generally true for both matrices: LLM finds higher levels of sensationalism and bias than humans.




Figure 4: Confusion Matrix for Criterion: Sensational Language. 1: Sensational; 4: Neutral




Figure 5: Confusion Matrix for Criterion: Article Bias. 1: Biased; 4: Unbiased

Figure 6 shows the gain in agreement when we use the refined prompt from Table 3 for the same two criteria, *ArtBias* and *SensLang*. The question remains unchanged, while the responses have become binary: Biased/Unbiased and Sensational/Neutral. What happens with the refined prompt is that k goes from 0.2064 to 0.4750 for *ArtBias* and from 0.1732 to 0.5486 for *SensLang*.




Figure 6: Comparison of agreement levels between the refined and initial implementations of the criteria, considering only *ArtBias* and *SensLang*.

5.2 Resolution of disagreements between human experts

In the previous section, we considered all the cases where human annotators agreed. We now focus on cases where annotators disagree, to see if LLM can help resolve disagreements. We perform an ex-post human evaluation to create a ground truth, and cases where even this additional analysis does not provide a definitive answer are considered borderline.

We focus on the *ArtBias*, *SensLang*, *NegTarg* and *HeadAcc* criteria because these are the ones where the annotators disagreed the most. Furthermore, we only consider cases of relevant disagreement between annotators, i.e., cases where one annotator answered yes and the other did not, or, in the case of multiple answers, there are at least two degrees of separation between their answers (e.g., one says Inaccurate and the other says Quite Accurate).

Table 4 shows, for the 5 criteria considered, (1) the total number of articles, (2) the number of articles where the human annotators gave different answers, (3) the number of times these disagreements are relevant, (4) the number of times the LLM gave the correct answer, and (5) the so-called borderline cases.

Articles where human annotators disagree on at least one criterion are 226. The table shows that, excluding borderline cases, the LLM resolves 100% of disagreements for *ArtBias* and *NegTarg*, 81% for *HeadAcc*, and 63% for *SensLang*. 100% for the two *NegTarg* cases comes from the fact that LLM always gives a correct answer on non-borderline cases. The same is true for *NegTarg* Identification.

Table 4: Statistics about disagreement between experts

criteriaID	no_articles / no_art_disagreement / relevant_disagreements / LLM-correct / Borderline
ArtBias	340 / 79 / 4 / 4 / 0
HeadAcc	340 / 108 / 11 / 9 / 0
NegTarg (Detection)	340 / 30 / 30 / 18 / 12
NegTarg (Identification)	340 / 47 / 47 / 32 / 15
SensLang	340 / 72 / 11 / 7 / 0

This highlights the LLM’s potential to effectively assist humans in resolving disagreements related to the evaluation of these criteria.

For the degree-based criterion *ArtBias*, we observe that the LLM resolves disagreements (i.e., its response coincides with the ex-post ground truth) by selecting the central evaluation among those provided by human annotators. For example, if one annotator selected Biased and the other annotator selected Quite Unbiased, the resolution is with the annotation Quite Biased.

For *NegTarg*, the LLM shows a strong ability to identify whether a group or individual is being negatively targeted within an article. However, there are several cases where even human experts find it difficult to make such determinations. The LLM shows similar proficiency when tasked with identifying the specific subject of negative targeting.

For *SensLang*, although the number of relevant disagreements is limited, as in the cases of *ArtBias*, we observe that the LLM is particularly sensitive to sensationalist content, especially content designed to evoke strong emotional responses in readers. For example, in nine cases where an annotator classified the content as "neutral", the LLM detected sensationalism levels later confirmed by the ground truth (corresponding to the ex-post human annotator).

6 Discussion

In this study, we investigated whether and to what extent one Large Language Model can automatically evaluate criteria that journalistic organizations use in real-world evaluations to estimate reliability of online publishers. We adopt a subset of criteria used by the Global Disinformation Index (and very similar to some used by NewsGuard) to assess the risk of an online publisher exposing its readers to misinformation. The first of the research questions, **RQ1**, is whether it is possible to define quality prompts to help LLMs correctly evaluate good journalism criteria to assess the reliability level of a publisher. The second of the research questions, **RQ2**, is how good is the LLM at aligning with the responses given by human experts, in the same context as before. To answer these questions, we implemented an iterative refinement of the prompts. This process allowed us to improve some of the initial prompt definitions for better processing by the LLM. As described in Section 5.1, among the six criteria analyzed, the evaluation of three of them was already in good agreement with that of human experts. Specifically, in recognizing the nature of a news story, in recognizing the presence of a lede, and in understanding whether a text is negatively targeting something or someone. When it came to recognizing sensational or biased language, the initial prompts resulted in very low levels of agreement between the humans and the LLM. Upon deeper analysis of these results, we observe that LLMs show increased sensitivity in detecting bias and sensationalism compared to human experts. In addition, most errors are confined to neighboring classes, which mitigates the severity of prediction inaccuracies. In particular, the use of refined prompts (offering binary responses instead of four options) significantly increases agreement, especially in detecting sensational language.

RQ3 asks what types of questions LLMs can answer correctly and thus support the human annotator in the publisher's evaluation process. To address this third question, we first note that in scenarios where increased annotation sensitivity is acceptable or a reduction in predictive granularity is warranted (e.g., considering binary instead of multiclass responses), the LLM-based

annotator provides good quality annotations. As anticipated, this is particularly evident for criteria requiring significant effort, such as detecting the use of sensational language or assessing article bias in the text. Even better when it comes to judging whether the article is negatively targeting someone or something (for this criterion, *NegTarg*, Detection, we got the highest k of 0.7089).

In addition, an LLM-based annotator can offer valuable assistance in resolving disagreements among human experts. Specifically, the LLM demonstrates promise in supporting decisions concerning *ArtBias*, *HeadAcc*, and *NegTarg*. This includes detecting negatively targeted groups or individuals and accurately identifying them—tasks that are often complex and may not lead to definitive conclusions. In particular, accurate assessment of negative targeting of individuals or minority groups is critical to combating the spread of adversarial narratives. [1]. The significant alignment achieved in a zero-shot experimental setting underscores the potential of the LLM for further improvement through tuning and adaptation.

7 Conclusions and Limitations

We explored the use of one LLM to evaluate a set of criteria designed to assess online journalism standards. The LLM achieves substantial levels of agreement with human experts on complex evaluations, such as understanding whether the text negatively targets someone or something. In addition, the LLM is highly effective at resolving disagreements among human experts in tasks such as detecting negative targeting and determining bias and headline accuracy in articles.

While our study provides valuable insights into the use of Large Language Models (LLMs) to support human judgment in evaluating news quality, it is subject to certain limitations. However, we believe that these limitations do not compromise the core objective of our work, which is to demonstrate that LLMs can be a promising tool for this task. A primary limitation concerns the scope of our dataset, which is relatively limited in terms of both size and linguistic coverage, as it only includes Italian news articles. An expansion of this dataset would allow a more comprehensive evaluation of the proposed methodology in different linguistic contexts. Furthermore, although we have considered a number of relevant reliability criteria, our evaluation does not cover all possible aspects of news quality assessment. For example, visual elements such as the likelihood of an image being automatically generated or having sensationalist characteristics were not included. Furthermore, while our experiments focused on a single LLM, GPT-4o, to demonstrate the feasibility of our approach, a broader investigation of the performance of different models - including both proprietary and open source alternatives of different sizes - would be beneficial. Moreover, alternative retrieval strategies, such as Retrieval-Augmented Generation (RAG) [28], one-shot and few-shots learning could further enhance the robustness and adaptability of the methodology.

Building on this foundation, future work will focus on extending the scope of our experiments in several key directions. First, we plan to extend the evaluation criteria to include visual aspects, thereby capturing a more holistic representation of news quality. Second, we aim to increase the number of news outlets and articles analyzed in order to improve the robustness and generalizability of our findings. In addition, we will explore the applicability of our

approach across multiple languages to evaluate its effectiveness in different linguistic environments. Finally, we intend to experiment with a wider range of LLMs and integrate advanced techniques such as RAG and few-shot learning to refine our evaluation framework and further improve performance.

Acknowledgments

This work is partially supported by SERICS (PE00000014) under the NRRP MUR program funded by the EU - #NGEU

References

- [1] 2022. Disinformation as Adversarial Narrative Conflict. <https://www.disinformationindex.org/blog/2022-06-22-disinformation-as-adversarial-narrative-conflict/>.
- [2] 2024. Ad Fontes Media. <https://adfontesmedia.com/>.
- [3] 2024. European Digital Media Observatory (EDMO) Project. <https://edmo.eu/resources/training-programme/>.
- [4] 2024. Facebook's Tips to Spot False News. <https://www.facebook.com/formedia/blog/third-party-fact-checking-tips-to-spot-false-news>.
- [5] 2024. The Global Disinformation Index. <https://www.disinformationindex.org/>.
- [6] 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>.
- [7] 2024. Iffy Index. <https://iffynews/index/>.
- [8] 2024. MediaBias Fact Check. <https://mediabiasfactcheck.com/>.
- [9] 2024. NewsGuard. <https://www.newsguardtech.com>.
- [10] 2024. OpenAI GPT-4o. <https://platform.openai.com/docs/models/gpt-4o>.
- [11] 2024. Scikit-learn cohen's kappa score implementation. https://scikit-learn.org/stable/modules/generated/sklearn.metrics.cohen_kappa_score.html.
- [12] 2024. Selenium. <https://www.selenium.dev/>.
- [13] 2024. The Trust Project - trust indicators. <https://thetrustproject.org/trust-indicators/>.
- [14] 2024. wget. <https://www.gnu.org/software/wget/>.
- [15] 2024. Writing differences feature, news and investigative articles. <https://gspieler.medium.com/whats-the-difference-writing-hard-news-feature-and-investigative-journalism-40811e0092c>.
- [16] Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational linguistics* 34, 4 (2008), 555–596.
- [17] John W. Ayers, Adam Poliak, Mark Dredze, Eric C. Leas, Zechariah Zhu, Jessica B. Kelley, Dennis J. Faix, Aaron M. Goodman, Christopher A. Longhurst, Michael Hogarth, and Davey M. Smith. 2023. Comparing Physician and Artificial Intelligence Chatbot Responses to Patient Questions Posted to a Public Social Media Forum. *JAMA Internal Medicine* 183, 6 (06 2023), 589–596. <https://doi.org/10.1001/jamainternmed.2023.1838>
- [18] Parisa Bazmi, Masoud Asadpour, and Azadeh Shakery. 2023. Multi-view co-attention network for fake news detection by modeling topic-specific user and news source credibility. *Information Processing & Management* 60, 1 (2023), 103146.
- [19] Md Momen Bhuiyan, Amy X Zhang, Connie Moon Sehat, and Tanushree Mitra. 2020. Investigating differences in crowdsourced news credibility assessment: Raters, tasks, and expert criteria. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–26.
- [20] John Bianchi, Manuel Pratelli, Marinella Petrocchi, and Fabio Pinelli. 2024. Evaluating Trustworthiness of Online News Publishers via Article Classification. In *Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing*. 671–678.
- [21] Tatiana Celadin, Valerio Capraro, Gordon Pennycook, David G Rand, et al. 2023. Displaying news source trustworthiness ratings reduces sharing intentions for false news posts. *Journal of Online Trust and Safety* 1, 5 (2023).
- [22] 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>
- [23] M. Freeze, M. Baumgartner, P. Bruno, et al. 2021. Fake Claims of Fake News: Political Misinformation, Warnings, and the Tainted Truth Effect. *Polit Behav* (2021), 1433–1465. <https://doi.org/10.1007/s11109-020-09597-3>
- [24] Hendrik Heuer and Elena L. Glassman. 2024. Reliability Criteria for News Websites. *ACM Trans. Comput.-Hum. Interact.* 31, 2, Article 21 (Jan. 2024), 33 pages. <https://doi.org/10.1145/3635147>
- [25] Maurice Jakesch, Jeffrey T. Hancock, and Mor Naaman. 2023. Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences* 120, 11 (2023), e2208839120. <https://doi.org/10.1073/pnas.2208839120> arXiv:<https://www.pnas.org/doi/pdf/10.1073/pnas.2208839120>
- [26] Antino Kim, Patricia L. Moravec, and Alan R. Dennis. 2019. Combating Fake News on Social Media with Source Ratings: The Effects of User and Expert Reputation Ratings. *Journal of Management Information Systems* 36, 3 (2019), 931–968. <https://doi.org/10.1080/07421222.2019.1628921>
- [27] Anastasia Kozyreva, Philipp Lorenz-Spreen, Stefan M Herzog, Ullrich KH Ecker, Stephan Lewandowsky, Ralph Hertwig, Ayesha Ali, Joe Bak-Coleman, Sarit Barzilai, Melisa Basol, et al. 2024. Toolbox of individual-level interventions against online misinformation. *Nature Human Behaviour* (2024), 1–9.
- [28] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [29] Hause Lin, Jana Lasser, Stephan Lewandowsky, Rocky Cole, Andrew Gully, David G Rand, and Gordon Pennycook. 2023. High level of correspondence across different news domain quality rating sets. *PNAS Nexus* 2, 9 (09 2023), pggad286. <https://doi.org/10.1093/pnasnexus/pgad286>
- [30] Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica* 22, 3 (2012), 276–282.
- [31] NewsGuardCriteria. 2024. NewsGuard's rating process criteria. <https://www.newsguardtech.com/ratings/rating-process-criteria/>.
- [32] NewsGuardCriteriaProduct [n.d.]. GDI product and FAQ. <https://www.disinformationindex.org/product/>.
- [33] NewsGuardSafetySuite. 2024. NewsGuard's AI Safety Suite. <https://www.newsguardtech.com/solutions/ai-safety-suite/#>.
- [34] Gordon Pennycook, Adam Bear, Evan T. Collins, and David G. Rand. 2020. The Implied Truth Effect: Attaching Warnings to a Subset of Fake News Headlines Increases Perceived Accuracy of Headlines Without Warnings. *Management Science* 66, 11 (2020), 4944–4957. <https://doi.org/10.1287/mnsc.2019.3478>
- [35] Marinella Petrocchi, Manuel Pratelli, and Angelo Spognardi. 2022. Disinformation Risk Assessment: The Online News Market in Italy. <https://www.disinformationindex.org/country-studies/2022-1-31-the-online-news-market-in-italy/>
- [36] Manuel Pratelli and Marinella Petrocchi. 2022. A Structured Analysis of Journalistic Evaluations for News Source Reliability. In *Workshop Proceedings of the 16th International AAAI Conference on Web and Social Media, ICWSM 2022 Workshops*, Pedro O. S. Vaz de Melo, Wei Jeng, and Cody Buntain (Eds.). <https://doi.org/10.36190/2022.51>
- [37] Manuel Pratelli, Fabio Saracco, and Marinella Petrocchi. 2024. Unveiling News Publishers Trustworthiness Through Social Interactions. In *Proceedings of the 16th ACM Web Science Conference (Stuttgart, Germany) (WEBSCI '24)*. Association for Computing Machinery, New York, NY, USA, 139–148. <https://doi.org/10.1145/3614419.3644015>
- [38] Toby Prike, Lucy H Butler, and Ullrich KH Ecker. 2024. Source-credibility information and social norms improve truth discernment and reduce engagement with misinformation online. *Scientific Reports* 14, 1 (2024), 6900.
- [39] Kai-Cheng Yang and Filippo Menczer. 2023. Large language models can rate news outlet credibility. *CoRR* abs/2304.00228 (2023). <https://doi.org/10.48550/ARXIV.2304.00228> arXiv:2304.00228
- [40] Junjie Ye, Xuanning Chen, Nuo Xu, Can Zu, Zekai Shao, Shichun Liu, Yuhan Cui, Zeyang Zhou, Chao Gong, Yang Shen, Jie Zhou, Siming Chen, Tao Gui, Qi Zhang, and Xuanjing Huang. 2023. A Comprehensive Capability Analysis of GPT-3 and GPT-3.5 Series Models. arXiv:2303.10420 [cs.CL]
- [41] Amy X Zhang, Aditya Ranganathan, Sarah Emlen Metz, Scott Appling, Connie Moon Sehat, Norman Gilmore, Nick B Adams, Emmanuel Vincent, Jennifer Lee, Martin Robbins, et al. 2018. A structured response to misinformation: Defining and annotating credibility indicators in news articles. In *Companion Proceedings of the The Web Conference 2018*. 603–612.

A Ethical Considerations

Our research involves the use of LLMs to evaluate the reliability of online news publishers, a task which includes a step of annotating news articles according to journalistic criteria, usually done by human annotators. While the ability of LLMs to assist in such evaluations offers some advantages, it is imperative to recognize and address the potential ethical implications of this technology.

- Bias and Discrimination: Given that LLMs are trained on large datasets, they may inherently carry biases present in their training materials. In the context of news evaluation, such biases could affect the impartiality of trustworthiness assessments, potentially perpetuating stereotypes or unfairly targeting certain publishers based on biased training data.
- Transparency and Accountability: The decision-making process of LLMs is typically complex and not fully transparent, making it difficult to understand how or why certain scores

are generated. This lack of clarity can be particularly problematic in scenarios where decisions need to be justified or contested.

- **Environmental Impact:** The computational demands of training and running LLMs are substantial, contributing to significant energy consumption and environmental impact. This is a critical consideration given the global push to reduce carbon footprints.

In using LLMs to support the evaluation of news publishers' reliability, it is important not only to harness their potential to improve efficiency and accuracy, but also to rigorously address these ethical concerns. By actively working to mitigate bias, increase transparency, and reduce environmental impact, we can better harness the power of LLMs in a way that is consistent with ethical standards and societal expectations.

B Online Resources

Datasets and code to reproduce our experiments can be found *here*.

C Additional data

Table 5: The dataset analyzed in this study comprises Italian online publishers, as sourced from [35].

Publisher	Website
Avvenire	www.avvenire.it
Il Corriere Del Giorno	www.ilcorrieredelgiorno.it
Il Piccolo	www.ilpiccolo.gelocal.it
Corriere Della Sera	www.corriere.it
Il Fatto Quotidiano	www.ilfattoquotidiano.it
Il Post	www.ilpost.it
Domani	www.editorialedomani.it
Il Foglio	www.ilfoglio.it
Il Primato Nazionale	www.ilprimatonazionale.it
Il Gazzettino	www.ilgazzettino.it
Il Giornale	www.ilgiornale.it
Il Quotidiano Del Molise	www.quotidianomolise.com
Il Giornale Di Sicilia	www.gds.it
Il Giorno	www.ilgiorno.it
Il Resto Del Carlino	www.ilrestodelcarlino.it
Il Manifesto	www.ilmanifesto.it
Il Secolo D'Italia	www.secoloditalia.it
Il Mattino	www.ilmattino.it
Il Messaggero	www.ilmessaggero.it
Il Sole 24 Ore	www.ilsole24ore.com
Il Tirreno	www.iltirreno.gelocal.it
La Gazzetta Del Mezzogiorno	www.lagazzettadelmezzogiorno.it
La Nuova Ferrara	www.lanuovaferrara.gelocal.it
La Nazione	www.lanazione.it
La Nuova Padania	www.lanuovapadania.it
La Nuova Sardegna	www.lanuovasardegna.it
La Repubblica	www.repubblica.it
La Stampa	www.lastampa.it
La Verità	www.laverita.info
Libero	www.liberoquotidiano.it
Libertà	www.liberta.it
L'Unione Sarda	www.unionesarda.it
Open	www.open.online
Stopcensura	www.stopcensura.online