

Why People Trust Wikipedia Articles: Credibility Assessment Strategies Used by Readers

Houda Elmimouni
helmimo@iu.edu

University of Indiana, Bloomington
Bloomington, USA

Andrea Forte
Drexel University
Philadelphia, PA, USA
af468@drexel.edu

Jonathan Morgan
CrowdStrike
USA
jonnymorgan.esq@gmail.com

ABSTRACT

We examine how a diverse global readership assigns trust to Wikipedia articles, and the strategies they use to assess Wikipedia’s credibility. Through surveys and interviews, we develop and refine a Wikipedia trust taxonomy that describes the mechanisms by which readers assess the credibility of Wikipedia articles. Our findings suggest that readers draw on direct experience, established online content credibility indicators, and their own mental models of Wikipedia’s editorial process in their credibility assessments. Our findings can help the development of general online information assessment frameworks and the design of open collaboration systems to support credibility evaluation and trust calibration.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

KEYWORDS

credibility assessment, trust, Wikipedia, trust taxonomy

ACM Reference Format:

Houda Elmimouni, Andrea Forte, and Jonathan Morgan. 2022. Why People Trust Wikipedia Articles: Credibility Assessment Strategies Used by Readers. In *The 18th International Symposium on Open Collaboration (OpenSym 2022)*, September 7–9, 2022, Madrid, Spain. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3555051.3555052>

1 INTRODUCTION

From the mid 2000s to early 2010s, novel forms of information production and access like news aggregators, wikis, and social media appeared poised to transform the ways people find and use information. Research on information credibility assessment in such environments bloomed and scholars anticipated a need for updated models and began investigating how people discovered and made sense of information in new online contexts. A change was clearly coming.

At the time of this writing in 2022, the transformation is well underway. Wikis, aggregators, and social media are no longer new

features of people’s online experience, they are mundane everyday sources of information that shape the public discourse and the behavior for billions of people around the world. Unexpected outcomes abound: disinformation campaigns, hoaxes, and rumors have prompted efforts to automate and extend fact-checking; yet, in the midst of it, nearly two decades after its launch, the “encyclopedia that anyone can edit” persists as the most accessed reference resource in history. How do our hard-won models of credibility assessment stand up?

To understand Wikipedia readers’ credibility assessment strategies, we conducted a three-stage study. In the first phase, we deployed a Wikimedia-hosted survey on English Wikipedia to ask open-ended questions to readers as they read Wikipedia articles in the course of everyday internet use. The survey was presented to random Wikipedia readers as they used the site. Survey questions prompted readers to briefly describe what they noticed about articles and how they assessed article quality. We then used these data to develop a “trust taxonomy” of credibility assessment strategies. In the next phase, we used the taxonomy to design a second survey to understand the prevalence of these different strategies with a larger sample of readers, and to examine how these strategies related to different aspects of a reader’s motivation and information context, such as their prior familiarity with the current topic and their motivation for reading that particular article at that time. This survey was also shown to a random selection of Wikipedia readers. Lastly, we conducted interviews with a subset of 17 respondents to the second survey and asked them to elaborate on their responses to collect richer data on credibility assessment of Wikipedia articles.

Our study makes several contributions to the domain of online credibility assessment: 1– it captures information assessment behaviors in a natural setting during the course of everyday Wikipedia use. We deployed two intercept surveys on English Wikipedia, asking readers to provide their responses in the context of the specific articles that respondents had chosen to read for their own reasons, not as part of an experiment or intervention. 2– The embedded surveys allowed us to capture a geographically and demographically diverse pool of Wikipedia readers, unlike many studies that are constrained by the local availability of paid research panel participants or students. 3– Our inductive approach to trust category development relies on few assumptions about how people assign trust to Wikipedia content is like or unlike other online information sources. Our bottom-up approach yields constructs that are grounded in practices related to Wikipedia. 4– Comparing our taxonomy with frameworks used in previous studies of Wikipedia trust, we discuss how our model identifies several bases for trust that may help explain some of the contingent and ambiguous findings from previous survey-based Wikipedia trust studies.

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

OpenSym 2022, September 7–9, 2022, Madrid, Spain

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9845-9/22/09...\$15.00

<https://doi.org/10.1145/3555051.3555052>

1.1 Research questions

Our three-phase study answers the following research questions:

- What strategies do Wikipedia readers use when assessing the credibility of Wikipedia articles?
- How are these assessment strategies related to trust in Wikipedia articles?
- What factors influence choice of assessment strategy?

2 LITERATURE REVIEW

2.1 Trust in online information

Researchers have been interested in understanding how and why people assess the credibility of information on the internet since at least the late 1990s [6]. Broadly speaking, many of these researchers have recognized a distinction between indicators of credibility based on the content of information sources (broadly speaking, the visual, structural, and textual elements of the web page), and indicators based on the context, including the reader's perceptions of the web site's purpose, the social context in which the topic is situated, and their own background and motivation for accessing the information. Findings from Fogg et al. [5] web survey, which elicited freeform in-context credibility assessments from over 2,000 participants across a variety of web sites, suggest that people rely most on explicit content cues such as visual design, information structure, and information focus, but that contextual cues, such as the individual's assessment of the website's motive and reputation, and their perceptions of the usefulness of the information presented, are also important factors [5]. In another large-scale web survey, Flanagin and Metzger [4] found that online credibility assessments were related to people's prior familiarity with the information provider, and that their likelihood of verifying the accuracy of online information was related to their motivation for seeking it out in the first place. The importance of contextual, rather than content-based, factors in online credibility judgements is supported by Hargittai et al. [13], who find that students tend to make credibility judgements based on factors such as how highly the information source is ranked in search engine result pages (SERPs). More recently, Zhang et al. [29] also found that in the case of online news articles, domain experts tend to weigh contextual factors such as the presence of ads or indicators of whether the article had been independently fact-checked higher than non-experts, than with factors related to the actual contents of the article.

2.2 Trust in Wikipedia

The degree to which people trust information on Wikipedia is likely influenced by many of the same content- and context-related factors that drive trust in other prominent online information sources. This supposition is supported by research. Forte et al. [7] demonstrated that using lightweight visualizations to increase the salience of markers of information provenance (a content-based factor) and information stewardship (a contextual factor) both affected rater's assessments of the overall quality of the article. Similarly, Kittur et al. [17] found that highlighting aspects of article stability supported article assessment. Rowley et al. [23] also identified a range of content and contextual factors in their students' assessments of Wikipedia quality.

However, Wikipedia is different from other online information sources in at least two ways that may influence reader trust. First, Wikipedia is an encyclopedia and provides the largest and most comprehensive source of curated textual information on the open web; it is not intended to be the authority on any specific topic. Second, Wikipedia follows an open collaboration-based contribution model, with no centralized editorial review [8, 9]. A third difference also bears mentioning: Wikipedia content generally ranks highly in search results, to the point where removing Wikipedia content from search results can measurably degrade peoples' perceptions of the usefulness of those results [20]. The Wikimedia Foundation's own research shows that in some countries readers are so accustomed to accessing Wikipedia content through Google that they may not recognize the difference between the two sites [10]. These findings suggest that because of Wikipedia's ubiquity among online information sources it may not be easy to separate the reputation of Wikipedia itself (and therefore trust in Wikipedia as a brand) with that of the search platforms that drive readers to Wikipedia in the first place. In our study, we attempted to mitigate this confounding factor by surveying people within the context of Wikipedia itself. We believe that asking people to make credibility assessments of Wikipedia while they were reading Wikipedia articles will increase the salience of Wikipedia's brand identity and reputation for our survey respondents.

Despite the concerns voiced by educators and scholars about the accuracy of Wikipedia's content [3], surveys that have asked how much people trust Wikipedia in general have almost universally demonstrated that overall trust is extremely high. A 2014 YouGov survey found that respondents in the UK trust Wikipedia more than almost any other media institution [15]. A survey conducted by YouGov in 2019 found consistently high trust in Wikipedia across five countries polled: between 78% and 98% of respondents said they trusted Wikipedia "A great deal" or "somewhat" [26].

Our current understanding of whether, how, and why people trust Wikipedia is limited by the overwhelming focus of the relevant research to date on student populations. Studies of Wikipedia in educational contexts show that students express a high degree of trust in Wikipedia across countries and education levels [22]. Among student users of Wikipedia, trust in Wikipedia can be mediated by their particular information-seeking task, their prior knowledge of the subject, the specific topic itself (with topics viewed as "educational" eliciting higher trust). Trust in Wikipedia may also be moderated by negative attitudes expressed by teachers, who have historically cautioned students against using Wikipedia in academic contexts [22]. Rowley et al. [23] provide a thorough review of this subject.

Research on what motivates people to turn to Wikipedia, as opposed to some other information source, has also focused on student populations or on specific subject domains, such as health information. For example, studies of Wikipedia in educational contexts show that students consistently rely on Wikipedia for information related to their work [21, 22]. But students are not the only people who use Wikipedia, and Wikipedia articles on a multitude of topics are used to meet a wide variety of information-needs unrelated to formal education. English Wikipedia alone currently receives over 500 million unique monthly visitors [28]. In order to understand why trust in Wikipedia is so high, and what credibility assessment

strategies readers deploy to assess the credibility of the information they find on Wikipedia, we need to understand why people in general read Wikipedia.

2.3 Why people read Wikipedia

For over a decade, Wikipedia has consistently ranked among the top 10 websites by traffic [27]. The largest and most representative surveys of why people read Wikipedia so far confirm that seeking information related to work or school is an important driver of Wikipedia readership, but not the primary one. Research conducted by Singer et al. [24] and Lemmerich et al. [18] in collaboration with the Wikimedia Foundation indicate that while between 10 and 30 percent of article traffic across 14 language editions of Wikipedia is driven by information needs related to work or school assignments, readers articulated many other nuanced motivations for reading Wikipedia. Other highly-cited motivations included information seeking related to current events (e.g. topic came up in conversation, references to the topic in the media), seeking information to make a personal decision, fun/boredom, and a personal intrinsic interest in the topic. The diversity of motivations among the general Wikipedia readership suggest the possibility that Wikipedia readers deploy a wider range of credibility assessment strategies than those previously observed among student populations.

The Singer and Lemmerich surveys, which used an on-site intercept survey approach similar to our own, demonstrate that Wikipedia readers are readily able to articulate not only why they read Wikipedia in general, but what information needs they hope to address by reading a particular Wikipedia article. They found that Wikipedia readers turned to Wikipedia in roughly equivalent degrees in order to gain an overview of a topic, an in-depth understanding of the subject matter, and to look up a quick fact. The prevalence of multiple levels of "information need" among readers can be considered a contextual variable that may influence their approach to assessing the credibility of the article at hand: e.g. a reader may trust an article to provide accurate information about a "quick fact" like a celebrity's birthday or the name of the author of a famous book, but not trust the article to provide a complete and unbiased summary of the subject. We believe that the level of trust one assigns to a piece of information—whether at the level of a specific fact or an extended description of a topic—is an important facet of reader motivation that deserves study. Given the previous research on how people use both content and contextual cues to assign credibility to online information, we should expect that Wikipedia readers employ a variety of empirical and heuristically-driven strategies to assess whether they should trust what they read on Wikipedia—in other words, whether they judge the information they read on Wikipedia to be sufficiently credible to address their specific, immediate information need.

3 STUDY DESIGN AND METHODS

We undertook a three-phase study to characterize the state of information credibility practices of English Wikipedia readers around the globe. We do not claim a globally or linguistically representative sample; however, we cast a broad net and surveyed 1716 people from 105 countries using a Wikimedia-hosted survey as they looked for information on English Wikipedia to understand their credibility

assessment practices. Then, we conducted interviews with a subset of these Wikipedia readers to delve deeper into their reasoning. This study was approved by the IRB at Drexel University.

3.1 Phase 1. Establishing an Initial Taxonomy of Assessment Strategies

In phase 1, we used a survey hosted through an agreement with the Wikimedia Foundation. The survey request popped up randomly to Wikipedia readers while they viewed an article, leading to an off-platform survey. It randomly sampled 1:500 article page loads and it ran on English Wikipedia from January 7th through January 9th 2019. The survey was randomized by page view, rather than by article. Because Wikipedia traffic is highly skewed towards a relatively small set high-traffic articles related to topics of general interest, our results reflect the range and relative proportion of credibility assessment strategies that readers use for those kinds of articles.

The survey questions are inspired by Singer et al. [24], we used 4-point Likert-style questions to measure trust in Wikipedia as a whole and the current article in particular, and to assess how familiar the reader was with the article's topic. Subsequent questions were open-ended in order to elicit readers credibility-assessment strategies in their own words. We asked them to describe, in a sentence or two, why they trust or distrust Wikipedia in general, and the current article in particular. The text of the first survey is available on [Meta.Wikimedia.org](https://meta.wikimedia.org)¹

We received 297 complete survey responses to the first survey (425 responses total). Of the 359 responses for which we have platform data², 133 (37%) came from the desktop Wikipedia site and 226 (63%) came from the mobile site (en.m.wikipedia.org). For those who clicked the survey prompt, completion rate was slightly higher for respondents on desktop (74%) than mobile (69%).

Two researchers independently open coded random samples of the free-text responses and some responses were double coded so the researchers could discuss differences in their coding. The researchers iteratively refined the emerging categories into a set of 18 "trust components" related to the context and content of the article itself (e.g. whether the reader thinks it is written in a professional way, whether they believe it is a "top hit" in search results related to the topic), as well as the people who wrote the article, and the article's topic. These trust components served both as structured prompts in the second survey (Phase 2) formed the basis of our credibility assessment taxonomy, which we subsequently refined through axial coding and thematic analysis of interviews with readers (Phase 3).

3.2 Phase 2: Measuring trust components and mediating factors

In phase 2, the 18 trust components were deployed as answer prompts in a Likert-style question matrix. We used the same sampling rate and intercept approach as the first survey. Participants were asked to rate their level of agreement with each statement. The

¹[https://meta.wikimedia.org/wiki/Research:Why_readers_trust_Wikipedia/Surveys#Round_1_\(EN\)](https://meta.wikimedia.org/wiki/Research:Why_readers_trust_Wikipedia/Surveys#Round_1_(EN))

²Some platform data was not captured, due to a logging error that affects some browsers

text of the second survey is available on Meta.Wikimedia.org³. The aim was to understand the relative salience of each trust component among a broader group of readers, and how these factors related to three aspects of the reader's information need and current context: their prior familiarity with the current topic, their motivation for reading this particular article, and their overall trust in the article. The survey ran on English Wikipedia from March 19th to March 22nd 2019.

1419 people took the survey (answered at least 1 question), 807 people completed the survey (made it to the final page), and 522 answered all of the questions. Because of data logging errors, we were only able to record which platform the respondent used (mobile vs desktop site) for 1092 respondents. Of these, 380 (35%) took the survey on desktop, and 712 (65%) took the survey from the mobile site. 55 of those who completed the survey indicated that they would be willing to be contacted for further research; these respondents formed our participant pool for Phase 3 of the study.

3.3 Phase 3: Investigating the role of assessment strategies in trust

We interviewed a strategic sample of respondents to the second survey (17), (M=13, F=4) from the pool of 55 volunteers and asked them to elaborate on their responses to better understand how they think about their credibility assessments. A detailed description and a summary of our participants demographics are available on Meta.Wikimedia.org⁴

Participants were asked if they want to discuss their survey answers or select another article they prefer to discuss. All participants but one selected the same article and three others discussed additional articles besides the article they considered during the survey. The interview questions were semi-structured and used the participants' responses from the survey as question prompts to get more in-depth insights about their answers. The aim is to gain a deeper understanding into the factors that mediate a reader's trust of Wikipedia content, including but not limited to citations (e.g. when, how, and why they use citations). Interviews were conducted in Fall 2019, remotely in English via the participant's preferred medium (Skype, hangout, zoom, WhatsApp) recorded, then transcribed using a transcription service.

Interview transcripts were analyzed using a thematic analysis method [2]. We identified themes by coding the data line-by-line [1]. Thematic analysis goes beyond identifying and counting occurrences of words or phrases to identifying implicit ideas [11]. The first author used Dedoose to code the data and discussed the findings with the second and third authors to facilitate a critical process and strengthen the conceptual integrity of the codes [19].

The interview findings were used to validate, contextualize, organize, and expand on the trust components developed from the surveys. Our final taxonomy of credibility assessment strategies (24 strategies in 8 subcategories and 3 categories) is presented in

³[https://meta.wikimedia.org/wiki/Research:Why_readers_trust_Wikipedia/Surveys#Round_2_\(EN\)](https://meta.wikimedia.org/wiki/Research:Why_readers_trust_Wikipedia/Surveys#Round_2_(EN))

⁴https://meta.wikimedia.org/wiki/Research:Why_readers_trust_Wikipedia/Participant_demographics

Figure 1: Information need, familiarity, and trust

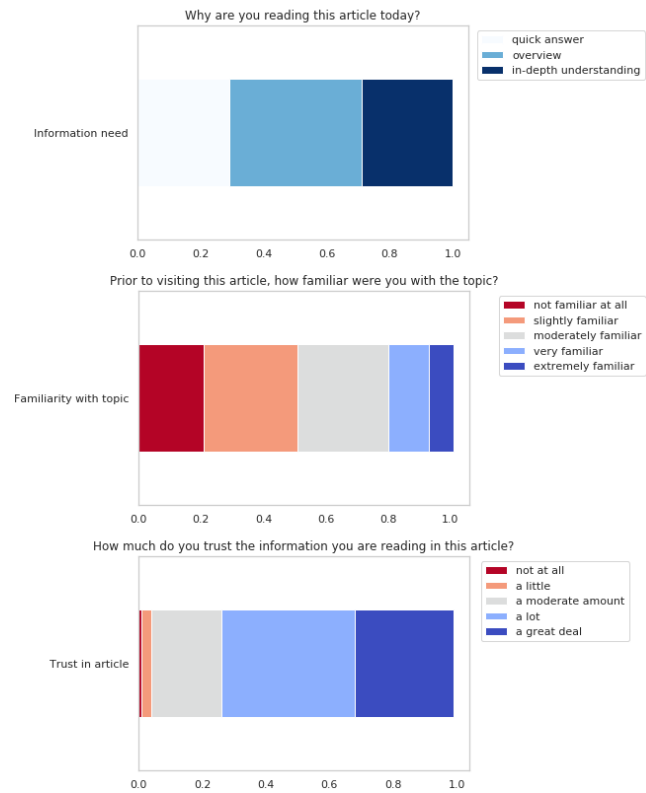


Figure 2. A full list with descriptions for each category and strategy can be found on Meta.Wikimedia.org⁵.

4 SURVEY FINDINGS: READER TRUST, FAMILIARITY, AND MOTIVATION

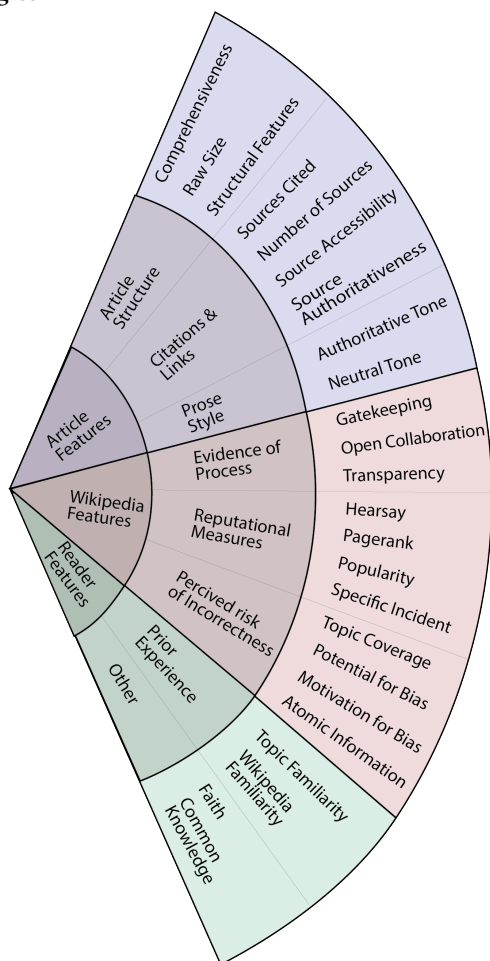
We first present and discuss descriptive statistics from our survey related to overall trust, familiarity, and motivation for reading. We then present and discuss findings related to readers responses to our trust component questions.

Overall, respondents reported a very high level of trust in Wikipedia. 88% of respondents to the first survey reported that they trusted Wikipedia as a whole "a lot" or "a great deal". 73% of respondents to the second survey reported that they trusted the information in the article they were currently reading "a lot" or "a great deal" (94% in the first survey⁶). In contrast, less than 4% of respondents in the second survey reported distrusting the information in the current article to any degree. This generally reflects the findings from the 2019 YouGov survey[26] as well as a reader survey commissioned

⁵https://meta.wikimedia.org/wiki/Research:Why_readers_trust_Wikipedia/Trust_taxonomy

⁶We believe that the 21% drop in current-article-trust between the first and second surveys is probably due to framing effects. In the first survey, the per-article trust question was directly preceded by the question "how much do you trust Wikipedia?" and a question about WHY they trusted Wikipedia as highly as they did. This framing may encourage people to evaluate the current article more positively than they would have otherwise. The second survey only asked about article-level trust and is probably more representative.

Figure 2: Taxonomy of Wikipedia Credibility-Assessment Strategies



by the Wikimedia Foundation in 2011 [16]. The relative proportion of responses to our "information needs" question are broadly similar to the results published in Singer et al. [24] and Lemmerich et al. [18] studies, which asked the same question to English Wikipedia readers a couple of years earlier. Familiarity with Wikipedia in the article's topic also generally aligns with Singer's results, with roughly half of respondents expressing at least "moderate" familiarity with the article's topic. Since the distribution of responses to these questions were consistent across both rounds of the survey, we only show results from the second survey. A breakdown of responses to all three questions is provided in **Figure 1**.

The 18 trust components presented in survey 2 (**Table 1**), which were developed based on the reasons for (dis)trusting Wikipedia elicited from respondents to the first survey, formed the first draft of our taxonomy. Participants were asked to evaluate the current article by indicating their level of agreement with each of these components. This allowed us to get a sense of which components are most salient to readers when it comes to making credibility assessments. The mean response value for answers to the trust component questions (from 3.7-4.4 out of 5) is fairly consistent overall, which suggest that these are, in fact, salient components of

credibility. The relatively high mean values, with higher agreement generally signifying a more positive assessment of the article's credibility, could also be interpreted as reflecting the high overall trust people place in Wikipedia. Analyzing the relationship between overall reported trust and responses to the trust component questions is beyond the scope of the current project, but would be a useful direction for future research.

Our findings show that Wikipedia readers consider a variety of contextual and content-based factors when making credibility judgements. The relative ranking of the trust components broadly align with the findings of Fogg et al. [5], Flanagan and Metzger [4] and Rowley et al. [23]. The four components that respondents find most salient (highest agreement) relate to the content of the article: assessments of the clarity and professionalism of the writing, the quality of the structure, and the accuracy of the information presented. The next four highest-ranked trust components focus on one aspect of the article's context, the characteristics of the article writers: their motivations (to present unbiased information, fix errors, help readers understand) and their perceived domain expertise. Intriguingly, readers do not seem to consider the "wisdom of the crowd" to be a particularly salient factor when making credibility assessments about Wikipedia articles: the three lowest-ranked trust components all relate, in one way or another, to the relationship between crowdsourcing and quality (search popularity, number of contributors, and number of reviewers). This finding suggests that, at least nowadays, reader trust in Wikipedia is not strongly influenced by either its status as one of the dwindling-number of prominent open collaboration platforms, or its ubiquity at the top of search results. Whether this is a consistent feature of Wikipedia's credibility, or an emergent property of changes within our online information ecology, is an intriguing subject for future research.

5 TAXONOMY: CREDIBILITY ASSESSMENT STRATEGIES OF WIKIPEDIA READERS

We present our final trust taxonomy and unpack the elements of the taxonomy in the context of subsequent interviews with a selection of Wikipedia readers recruited via the second survey. The taxonomy of 24 features that affect credibility assessment of Wikipedia articles, which we developed based on our reader surveys and interviews, is illustrated in **Figure 2**. In the sections below, we present the taxonomy's three categories, eight subcategories and their associated assessment strategies alongside quotes and observations gleaned from interviews with a selection of respondents from the second survey.

We encourage readers to take the following points into consideration when interpreting the taxonomy:

- Assessment strategies can be multi-valent: one person might articulate a specific strategy to explain why they trust an article, while someone else called out the same strategy to explain their distrust.
- Assessment strategies are not mutually exclusive: many interview participants drew on multiple strategies for assessing a single the article.
- Assessment strategies are not equally represented within the source data: as with the trust components, some strategies were more frequently articulated than others.

Table 1: Trust components from the second survey

Trust component	mean agreement	rank
<i>I believe that this article...</i>		
contains accurate information	4.34	3
contains an adequate number of references to external sources	4.06	12
contains detailed and comprehensive information	4.1	9
contains references to high quality external sources	3.97	13
has been read by many people	4.1	10
has been reviewed and corrected by many people	3.74	17
has been written by many people	3.73	18
is often a 'top hit' in search results related to the article topic	3.86	16
is well structured	4.37	2
is written in a professional way	4.31	4
is written in a way that is clear and easy to understand	4.43	1
<i>I believe that the people who write this article...</i>		
know a lot about the article topic	4.13	8
try to fix incorrect information when they see it	4.21	6
try to keep incorrect information from being added to the article	4.1	11
want the article to be neutral and unbiased	4.23	5
want to help readers understand how much to trust the information in the article	4.16	7
<i>I believe that the topic of this article...</i>		
has been written about in many other information sources (not just Wikipedia)	3.97	14
is written about in a neutral and unbiased way in other sources I have read	3.93	15

- Assessment strategies reflect the researchers' own understanding: our taxonomy is based on researchers' interpretations of statements made by interview and survey participants, which can be ambiguous.

5.1 Category 1: Reader Characteristics

These strategies reflect aspects of trust that stem from characteristics of readers themselves.

5.1.1 Prior Experience. This kind of assessment is based on the reader's personal experience with the content of Wikipedia. Such experience can be either a direct familiarity or Wikipedia familiarity.

- **Topic familiarity:** Degree that information the reader was looking for, or other information in the article, matches their prior knowledge of the subject
- **Wikipedia familiarity:** Degree that information on Wikipedia in general matches reader's prior knowledge

Many participants mentioned their prior experience as a reason for trusting Wikipedia. P9 mentioned multiple anecdotes about things she saw with her own eyes and that were accurate in Wikipedia. One time she mentioned: "I was born in Nepal. I was there as a kid, for quite a time. I came into the India ... once I was just reading an article in Wikipedia about something that is related to me, as I have seen it with my eyes...critical facts were written in that article.". P7 describes how his direct familiarity with the content of an article about the albums of a singer contributed to his trust of the content: "I was looking at a subject that I know a fair amount about. As I was looking through the list, there was nothing that seemed the slightest bit out of place. ... I'm familiar with every single one of the albums.". P11 summarizes why he trusts Wikipedia, encompassing both direct familiarity and Wikipedia familiarity: "So, my trust in the information in the article comes from multiple levels. One, I have used Wikipedia for many years as a resource, and I've grown to trust it. Two, ...And three, as I was reading the article, I had some basic knowledge of the university and universities in the country overall, and so all the information looked reasonable and believable. So, all these factors together combined to instill in me the trust that the information I was looking is accurate."

5.1.2 Other measures. While many assessments have justifiable basis, some other assessments have no clear or specific criteria.

- **Faith:** Unquestioned belief in the trustworthiness of the content without supporting rationale
- **Common knowledge:** Belief that the information in this article is widely or universally known and accepted

Some participants in our study expressed their full support to Wikipedia for reasons they could not explain. P4 illustrates such faith in Wikipedia: "For whatever reason, I support Wikipedia a lot. I'm not quite sure why frankly. I actually trust them more. I don't know why. I'm not always very trustworthy when I read things that people post or even in formal way. But definitely I do not trust a lot in the U.S. Media, particularly. But Wikipedia, I do trust them more. I don't know why. P15 uses his common sense to assess the article content. He explains: "there's just nothing here that's controversial enough to make me think, "Oh, well that can't be right." Or, something like that."

5.2 Category 2: Features of Wikipedia

These strategies reflect aspects of trust that stem from perceptions and features of Wikipedia as a whole.

5.2.1 Wikipedia process. Evidence or knowledge of how Wikipedia content is created contributes to people's trust in Wikipedia content. Process features can be related to open collaboration, evidence of gatekeeping, and transparency.

- **Open collaboration:** Low technical barriers to contribution and voluntary participation.
- **Evidence of gatekeeping:** Observations of specific indications that the article is actively monitored and moderated by people with decision-making authority
- **Transparency:** Ability to inspect the article development history and contributions

Although participants who never contributed to Wikipedia mentioned they do not know the details of how Wikipedia works, they have general ideas. P4 illustrates how a limited understanding of how Wikipedia works can lead to mistrust: "I was just thinking the other day, who posts Wikipedia? At one point I kind of like heard that

anybody can post a Wikipedia article or information about something. Is that true?" Those who contribute to Wikipedia, however, know more about the Wikipedia process. P14 described the collaborative aspect of Wikipedia as he compared it with other sources: "Wikipedia is much more collaborative, I sense that the people who go there are not trying to mislead people, but they're just trying to present information." For many participants, the fact that there is a kind of shepherding is an indicator of trustworthiness. P15 mentioned the case of controversial topic and how he senses gatekeeping while indicating the voluntary nature of Wikipedia writers: "Nobody gets paid for it. If there is like a very controversial topic then it's sometimes protected by Wikipedia. Not everyone can edit it unless they have like a good history of editing." Other participants mentioned citation needed notification or this article is incomplete notification. P10 describes the importance of such notifications as evidence of a quality control process: "I'm also spotting a citation needed, which is a good sign. I mean, obviously it's even better if there's a citation, but having a citation needed means somebody has looked at it and raised a query." A few participants did not know about the article development history but the participants who knew about it mentioned that they check the history to see who is editing, their number and what changes were made over time. For instance, P5 mentions how he checks the history: ". I look at the change history sometimes, just to say, 'Okay, what have they added that I maybe don't know about yet?' Or, if you find some exciting research and then you're like, 'Okay, I'll come back to this,' and then, you come back. You just managed to find it four years later. It's like, what changed?." Similarly P10 checks on the contributors: "Generally I look at history for that. I just go and read a section of the article and then scroll down and see who's writing what I would say, yeah, there's a decent amount of people working on this article from November 2012 I think."

5.2.2 Reputational measures. Some of the assessments are based on the perceived reputation of Wikipedia. This can be either based on the perceived popularity, the article PageRank, hearsay or a specific incident.

- **Popularity:** Perceptions about how many people consume or contribute to this article
- **PageRank:** Observation of the ranking of this article in search engine results pages
- **Hearsay:** General perceptions about how much other people trust this article, articles on this topic, or Wikipedia as a whole
- **Specific incident:** Indirect knowledge of specific incident(s) that influence credibility judgements

Participants expressed their perceptions of the popularity of the Wikipedia either in relation to the number of its users or contributors. P11 trusted the article he was reading based on his perception of users number. He says: "Just because I know ... very many people use Wikipedia as a resource. And I know the college probably has thousands of students, and alumni, and prospective students, and just curious people that might stumble upon it like myself. So to me it seems like a reasonable proposition that a lot of people have clicked on this article at some point or another." For P11, the number of users entails that if there was an error or missing information it would have been spotted and addressed quickly. P7 expressed his perception of the number of Wikipedia contributors: "I'll bet there's

a lot of writers that contribute to and edit this page. I would think that there's a strong chance that Donna Tartt herself might have made some contributions to this page." Some of the participants related their experience searching for Wikipedia articles in search engines and described how it ranked higher in the search results. P11 says: "... based on my own experience using the internet, that usually Wikipedia is among my top links. And often times, very often when I'm doing a first look for certain information, I'll do something. I'll type ... college Wiki in my search bar because I want to go right to Wikipedia." Furthermore, some participants mentioned how trust is based on other people trusting and recommending Wikipedia. P6 tells the story of how trust in Wikipedia has changed over time: "So, it's kind of funny, because I remember growing up, a lot of people, all the teachers in middle school, high school, were like, 'Don't use Wikipedia. Don't trust it. Don't trust it.' And, I found that started to change, especially in college where the teachers are looking at Wikipedia. It's like, 'All right. Oh, here you go. If you just want some more information about it, just go over there. It's right there. It's good.'" Last but not least, some participants related about incidents they experienced with Wikipedia that made them trust it more. Like P9 telling her story of what he knows about Napal, P16 and what he knew about Edward Hardwicke, and P7 about the Joni the musician. P6 told the story of how he watched an article developing and improving over time as he looked at the history interface: "I actually did that the other day for the first time... Then people would question. 'Is that the correct wording?' Or, 'Are you sure about that? Can you give me another source on that kind of thing?' So, I saw that. I wouldn't say I'm an expert on that part. Mostly because when I look in the change history, it's a lot of bots undoing stupid stuff."

5.2.3 Risk of incorrectness. Assessments based on the reader's judgement of the likelihood that this information could be wrong or misleading. These assessments can be based on topic coverage, atomic information, motivation for bias or potential for bias.

- **Topic coverage:** Perceived availability of reliable information on the topic in external sources
- **Atomic information:** Degree to which the information sought is simple or unambiguous
- **Motivation for bias:** Perceived likelihood that an unknown author would want to present wrong or misleading information on the topic
- **Potential for bias:** Degree to which information on this topic could be presented in a wrong or misleading way

Participants assessed the coverage of articles. In particular, the participants who were considering articles of contentious and controversial topics were sensitized to its coverage to other vantage point. For instance P2 mentions how every company page needs to have a controversy section to have the full coverage: "I think the bias inherent in every editor. I mean, that's the human nature I guess. But I think in general, yes, article writers tend to write in a balance no matter if not in the content directly in the structure of the content. So for example, ..., there has to be a controversy section for companies." P2's argument is very linked to the fact that while writing about companies there is a motivation for bias. P16 explains how the trust of an article can be affected when there is a motivation for bias in a political article: "Say compared to the impeachment that's going on right now. There's people have very strong opinions about that

sort of thing. So you want to make sure that you had, if you were crazy enough to discuss it with anyone, you'd want to make sure you have your facts straight. So it might be something I would read very closely. I would maybe read two or three times." Many other participants highlighted the potential for bias in political, religious and contentious topics in general. P5 describes such instance when he finds it hard to trust articles due to controversy: "If there's just some huge controversy or just debate where people get super pissed and angry about it, rather than just actually having an actual debate and discussion, yeah, I generally don't try to trust it as much." On the other hand some articles inherently have no potential for bias. P3 explains such bias potential in an article about bread: "It's just an article about a thing, about this type of bread. So I don't know how it could be biased in the first place. And just reading it, it's just about how it can be and how it's made and the bacteria in it. So, I guess this article particularly inherently, it's hard to make it biased I think." Presenting information in an unambiguous language makes the article more trustworthy. P16 details how the information presented made him rank an article about an actor as very trustworthy: "Absolutely. There's no obscure language in here. He's an actor, they're not using any sort of movie terms or anything that an average person wouldn't know. It's well written. It's clear cut and it's easy to comprehend."

5.3 Category 3: Article Features

Some credibility assessment strategies hinged on aspects of trust that stem from characteristics of the specific article being read.

5.3.1 Article structure. The structure of the article is a predominant criteria. It is based on the coverage/ comprehensiveness of the article, its overall size, or its structural features.

- **Perceived comprehensiveness:** Degree to which the article presents all relevant information on the topic.
- **Raw size:** The length of the article
- **Structural features:** Visual or organizational elements of the article content or user interface

Some participants assessed the trustworthiness of the article they are reading based on its perceived comprehensiveness, for example the existence of multiple perspectives or controversy. P2 mentioned that the absence a controversy section can highlight the absence of critical content: "I think it's missing, I don't know, a controversy section. There's not much critique about it. And I like that the most about Wikipedia articles and especially about companies is that they have a section about controversy." In addition to the coverage of the article, its length was also considered by the participants. Although, how long an article should be is relative, participants expected a substantial amount of information. P4 explained his stance vis-a-vis the article length: "I'd probably quantify like against myself for that matter, how much I read in there. I mean, you're right, there's nobody can say what is a lot. It depends on people, what people think is a lot or not a lot. But it just felt like there was a long information in there that I could read and learn so much from that article." While an article can be long and cover different aspects of the topic, organizational elements were also sought out and contributed to its assessment for trust. P2 explains the structural elements he considered: "I look at the structure of the article. I basically look at everything in the article. I look at the overall structure and by structure, I mean, the info box on the right and the headings. Does it have decent amount of headings

... but the biggest thing is probably the info box, the headings and references ... yeah, those four aspects I would say are my criteria for comparing the article to other articles on Wikipedia and judging if it's on par."

5.3.2 Citations and external links. These assessments are based on the prevalence or characteristics of cited sources or other external links. Participants checked whether the article has citations, how many, the reputation of the sources as well as if they can check and verify the source.

- **Presence of sources:** Whether the article contains (any) sources
- **Number of sources:** How many sources the article has
- **Perceived authoritativeness of sources:** the reputation or ethos of cited sources
- **Accessibility of sources:** The degree to which the information in the article may be independently verified by checking the cited sources

All participants highlighted the importance of citations in weighing the credibility of an article. For example, for P5 if an article is vague but has a citation he can trust it. He says "if it's something super vague and a little more simple, I guess, but it's still cited, fair enough. But otherwise, it can seem like someone's kind of just going off on a tangent... But I mean, overall, I mostly just look for the citations". The number of citations also matters in assessing the credibility of articles. P1 mentions the reasons why he assessed an article as very trustworthy: "Well, because it's presented from Wikipedia first things and second because there is over 50 different sources for this extremely well researched article." Participants also peruse the provenance of the citations and, in particular, their reputation. For example P10 moderately trusts the article as he cannot qualify the reputation of the sources. He says: "There are a couple of references I can see, two places I recognize like Reuters and the Chinese news agency, Xinhua, but most of them are to very obscure websites." Many of the participants mentioned how they appreciate being able to dig further and check the references, P2 describes how he avoids news articles as they may not be accessible and therefore of a questionable reliability: "ideally when I evaluate references, I try to avoid news articles a lot because my experiences with online news is that the information.. tends to change and editors go back and then change things. So most of the time they make a note of it and then other times the article disappears completely for whatever reason. Either it's redacted or just gone from the the archives... I guess it's a good resource to capture information and preserve information, but also raises the question of whether or not that information is still reliable."

5.3.3 Prose style. The prose style of the article is a popular indicator of the article credibility. Participants assessed the article based on the textual characteristics or writing style of the article.

- **Authoritative tone:** Degree to which the tone of the article is professional or suggests expertise
- **Neutral tone:** Degree to which the article contains biased or opinion-based language

P6 gives an illustration of a biased tone that would make her skeptical of the article: "If it's written like, "Oh, this policy's very good and very beneficial," like I mean, that's very vague. I don't know. If there was like, "This was an economic stimulant to the economy

which provided a great number of job," you know, it sounds a little more researched kind of thing, I guess." P11 expresses how he trusts an article due to the used tone: *"One other thing is also, with this article in particular, and many, most, if not all articles, it's written in straightforward, factual language. It's not written as opinion."*

6 DISCUSSION & IMPLICATIONS

The implications of our work are tri-fold:

- (1) Our taxonomy can be used to assess credibility in other online information sources that use a peer production model that are like Wikipedia. The taxonomy can be adapted as needed.
- (2) Our inductive, iterative, and contextual approach can be used to highlight gaps in other credibility assessment constructs or to develop tools to support credibility assessment and trust calibration for other general online sources.
- (3) Our taxonomy can inform the design of open collaboration systems to support credibility evaluation and trust calibration.

6.1 Exploratory analysis

The primary purpose of this study is to understand how much readers trust Wikipedia articles, and to develop a framework for characterising the content-based and contextual strategies they use when assessing the credibility of information in Wikipedia articles. However, the richness of our survey data-set provides an opportunity to contextualize these findings with supplementary statistical analysis. We believe that our study is the first to gather data related to trust in Wikipedia, motivations for reading, and topic familiarity from large and geographically diverse set of Wikipedia readers in context—while they were actually visiting Wikipedia to address their own information needs. In order to spark discussion and highlight potential areas of future research, we now present several exploratory statistical analysis of data from Survey 2.

6.1.1 Influence of information need and topic familiarity. We analyzed the relationship between readers' trust in a Wikipedia article and their information need. Our analysis shows that trust is significantly related to information need (Kruskal-Wallis $H=10.511$, $n=1350$, $p=0.005$). If "information need" is treated as an ordinal scale (with "quick fact" as the smallest information need and "in-depth understanding as the largest), there is a weak positive relationship between increased trust and size of information need (Spearman's Rho 0.078, $n=1350$, $p = 0.003$). Although the correlation is weak, this finding could indicate that readers have a higher threshold for trust when they require an in depth understanding of an article's topic vs. learning a quick fact contained within the article. Future work should investigate whether people with different information needs use different credibility assessment criteria, and whether different design interventions like those evaluated by Forte et al. [7] could help readers needs make better-informed judgements about whether a particular Wikipedia article (or other online resource) is likely to provide accurate information about what they came there to learn.

We also analyzed whether respondent who indicated that they were more familiar with the topic of the article tended to trust it

more, regardless of their information need. We did not find a significant correlation between self-professed trust in the information presented in the article and prior familiarity with the article's topic (Spearman's Rho 0.049, $n=1381$, $p = 0.07$). Although we found no correlation here, we posit that understanding topic familiarity—a proxy measure for domain expertise or background knowledge—may be a useful factor to include in future analysis of what mediates trust in online information resources.

6.1.2 Influence of article quality on trust. Our study was not aimed at determining whether readers are able to differentiate between relatively higher- and lower-quality Wikipedia articles. However, prior research has identified that overall perceptions of information quality influence trust in online resources, and that specific content features factor into credibility assessments. Therefore, we decided to investigate this potential relationship. Specifically, we tested whether there was a correlation between readers' trust in an article and that article's Wikipedia quality class, a six-level ordinal scale that ranges from low-quality "Stub" articles to high quality "Featured" articles. Because quality classes are determined based on the content of an article, and our respondents ranked content-focused measures as highly salient for determining credibility, we expect that higher quality articles will elicit higher trust overall. To perform this analysis, we used the ORES article quality model,[12] which uses structural features of Wikipedia articles to predict their quality class [25]. We found a (weak) positive relationship between a respondent's trust in an article and the predicted ORES quality class of that article (Spearman's Rho 0.067, $n=1312$, $p = 0.014$). This provides additional evidence that readers are able to accurately assess the general quality of the article they are reading, and that content-related factors do inform their credibility assessments.

6.1.3 Geographic differences in trust and information need. The second survey was taken by respondents from 105 different countries. We received at least 20 responses from 6 countries: USA, India, UK, Canada, Australia, and Germany. Below we present a geographic breakdown of responses, and results from an analysis of relationships between trust and information need across these countries. In line with previous trust questionnaires by Wikimedia [16], Jemielniak and Wilamowski [14], and YouGov [26], trust in Wikipedia varied by country or culture (Kruskal-Wallis $H=12.827$, $n=977$, $p=0.025$). On average, trust was highest among respondents in India and Germany and lowest in Canada and Australia, although a large variability in sample size between countries suggests caution in over-interpreting these results. Our data also revealed that respondents from different countries reported different patterns of information need (Chi $X=29.504$, $p=0.001$, $df=10$). We believe that these country-mediated relationships to Wikipedia deserve further investigation. A better understanding of the broad material (e.g. internet penetration, literacy rates) and cultural (e.g. epistemological bases for credibility, attitudes towards expertise) factors that influence why people use online information resources could help platforms serve their readers needs better. Additionally, further research could investigate whether these patterns exist across Wikipedia languages beyond English.

7 CONCLUSION

Our goal was to examine the strategies Wikipedia readership use to assess Wikipedia's credibility. We conducted a three-stage study where we surveyed a total of 1716 people from 105 countries and interviewed 17. We analyzed the data quantitatively and qualitatively, inductively and iteratively developed and refined a Wikipedia trust taxonomy. We found that Wikipedia readers draw on direct experience, established online content credibility indicators, and their own mental models of Wikipedia's editorial process in their credibility assessments. Our final taxonomy of credibility assessment strategies (24 strategies in 8 categories). We suggest our work can help the development of general online information assessment frameworks and the design of open collaboration systems to support credibility evaluation and trust calibration. Our study uncovered many areas for future work. It would be interesting to examine the relationship between overall reported trust and responses to the trust component questions, study how understanding credibility assessment strategies could help us understand and address misinformation, investigate whether people with different information needs or cultural background use different credibility assessment criteria, and whether different design interventions could help readers' needs make better-informed judgements about whether a particular Wikipedia article (or other online resource) is likely to provide accurate information about what they came there to learn.

ACKNOWLEDGMENTS

This work was partially supported by NSF grant #1253302. We also thank our interview and survey participants.

REFERENCES

- [1] Richard E Boyatzis. 1998. *Transforming qualitative information: Thematic analysis and code development*. sage.
- [2] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [3] Hsin-liang Chen. 2010. The perspectives of higher education faculty on Wikipedia. *The Electronic Library* (2010).
- [4] Andrew J. Flanagin and Miriam J. Metzger. 2000. Perceptions of Internet Information Credibility. *Journalism & Mass Communication Quarterly* 77, 3 (2000), 515–540. <https://doi.org/10.1177/107769900007700304> arXiv:<https://doi.org/10.1177/107769900007700304>
- [5] B J Fogg, Cathy Soohoo, David R Danielson, Leslie Marable, Julianne Stanford, and Ellen R Tauber. 2003. How Do Users Evaluate the Credibility of Web Sites?: A Study with over 2,500 Participants. In *Proceedings of the 2003 Conference on Designing for User Experiences (DUX '03)*. ACM, New York, NY, USA, 1–15. <https://doi.org/10.1145/997078.997097>
- [6] Brian J Fogg and Hsiang Tseng. 1999. The elements of computer credibility. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 80–87.
- [7] Andrea Forte, Nazanin Andalibi, Thomas Park, and Heather Willever-Farr. 2014. Designing information savvy societies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*. ACM Press, New York, New York, USA, 2471–2480. <https://doi.org/10.1145/2556288.2557072>
- [8] Andrea Forte and Amy Bruckman. 2008. Scaling consensus: Increasing decentralization in Wikipedia governance. In *Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)*. IEEE, 157–157.
- [9] Andrea Forte and Cliff Lampe. 2013. Defining, understanding, and supporting open collaboration: Lessons from the literature. *American behavioral scientist* 57, 5 (2013), 535–547.
- [10] Wikimedia Foundation. 2016. *New Readers: research and findings*. Technical Report. Wikimedia Foundation. https://meta.wikimedia.org/wiki/File:New_Readers_research_findings_presentation,_September_2016.pdf
- [11] Greg Guest, Kathleen M MacQueen, and Emily E Namey. 2011. *Applied thematic analysis*. Sage Publications.
- [12] Aaron Halfaker and R. Stuart Geiger. 2020. ORES: Lowering Barriers with Participatory Machine Learning in Wikipedia. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW2, Article 148 (oct 2020), 37 pages. <https://doi.org/10.1145/3415219>
- [13] Eszter Hargittai, Lindsay Fullerton, Ericka Menchen-Trevino, and Kristin Yates Thomas. 2010. Trust online: Young adults' evaluation of web content. *International journal of communication* 4 (2010), 27.
- [14] Dariusz Jemielniak and Maciej Wilamowski. 2017. Cultural diversity of quality of information on Wikipedias. *Journal of the association for information science and technology* 68, 10 (2017), 2460–2470.
- [15] William Jordan. 2014. British people trust Wikipedia more than the news. *Online document*. Accessed 11 (2014).
- [16] Ayush Khanna and Mani Pande. 2011. *Readers in US, Russia, Germany and India are the most pleased with Wikipedia Article Quality*. Technical Report. Wikimedia Foundation. <https://diff.wikimedia.org/2011/11/09/wikipedia-quality/>
- [17] Aniket Kittur, Bongwon Suh, and Ed H Chi. 2008. Can you ever trust a wiki?. In *Proceedings of the ACM 2008 conference on Computer supported cooperative work - CSCW '08*. ACM Press, New York, New York, USA, 477. <https://doi.org/10.1145/1460563.1460639>
- [18] Florian Lemmerich, Diego Sáez-Trumper, Robert West, and Leila Zia. 2019. Why the World Reads Wikipedia. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining - WSDM '19*. ACM Press, New York, New York, USA, 618–626. <https://doi.org/10.1145/3289600.3291021> arXiv:1812.00474
- [19] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and Inter-Rater Reliability in Qualitative Research: Norms and Guidelines for CSCW and HCI Practice. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 72 (Nov. 2019), 23 pages. <https://doi.org/10.1145/3359174>
- [20] Connor McMahon, Isaac Johnson, and Brent Hecht. 2017. The substantial interdependence of Wikipedia and Google: A case study on the relationship between peer production communities and information technologies. In *Eleventh international AAAI conference on web and social media*.
- [21] Jonathan T. Morgan and Mark Zachry. 2016. *Student use of free online information resources*. Technical Report. Wikimedia Foundation, University of Washington. https://meta.wikimedia.org/wiki/Research:Student_use_of_free_online_information_resources
- [22] Josiane Mothe and Gilles Sahut. 2018. How trust in Wikipedia evolves: a survey of students aged 11 to 25. *Information Research: an international electronic journal* 23, 1 (2018), n1.
- [23] Jennifer Rowley, Frances Johnson, and Laura Sbaffi. 2017. Gender as an influencer of online health information-seeking and evaluation behavior. *Journal of the Association for Information Science and Technology* 68, 1 (jan 2017), 36–47. <https://doi.org/10.1002/asi.23597>
- [24] Philipp Singer, Florian Lemmerich, Robert West, Leila Zia, Ellery Wulczyn, Markus Strohmaier, and Jure Leskovec. 2017. Why We Read Wikipedia. In *Proceedings of the 26th International Conference on World Wide Web - WWW '17*. ACM Press, New York, New York, USA, 1591–1600. <https://doi.org/10.1145/3038912.3052716> arXiv:1702.05379
- [25] Morten Warncke-Wang, Dan Cosley, and John Riedl. 2013. Tell Me More: An Actionable Quality Model for Wikipedia. In *Proceedings of the 9th International Symposium on Open Collaboration (Hong Kong, China) (WikiSym '13)*. Association for Computing Machinery, New York, NY, USA, Article 8, 10 pages. <https://doi.org/10.1145/2491055.2491063>
- [26] YouGov Wikimedia Foundation. 2019. *Women and Wikipedia: Summary of results from 2019 YouGov survey*. Technical Report. Wikimedia Foundation. https://commons.wikimedia.org/wiki/File:Women_%26_Wikipedia_-_Summary_of_results_from_2019_YouGov_survey.pdf
- [27] Wikipedia. [n.d.]. List of most visited websites. Accessed June 12, 2022. https://en.wikipedia.org/wiki/List_of_most_visited_websites
- [28] Wikipedia. [n.d.]. Wikipedia Monthly Overview Statistics. Accessed June 19, 2022. <https://stats.wikimedia.org/#/en.wikipedia.org>
- [29] Amy X Zhang, Martin Robbins, Ed Bice, Sandro Hawke, David Karger, An Xiao Mina, Aditya Ranganathan, Sarah Emlen Metz, Scott Appling, Connie Moon Sehat, Norman Gilmore, Nick B Adams, Emmanuel Vincent, and Jennifer Lee. 2018. A Structured Response to Misinformation. In *Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18*. ACM Press, New York, New York, USA, 603–612. <https://doi.org/10.1145/3184558.3188731>