

Towards a Unified Framework for the UX Design of News Credibility Tools

James Stomber
Credibility Coalition
jnmstomber@gmail.com

Dilrukshi Gamage
Tokyo Institute of Technology
dilrukshi.gamage@gmail.com

Bill Skeet
Credibility Coalition
bill@skeet.com

Amy X. Zhang
University of Washington
axz@cs.uw.edu

Abstract

In light of the proliferation of mis- and disinformation on social media platforms, a number of tools and technological interventions have been developed that attempt to detect, verify, or evaluate digital news media at scale. Such ‘credibility tools’ provide labels or indicators that help users identify credible content and avoid deceptive media. Little is known about the efficacy of such tools or how they are perceived by users; moreover, it is unclear whether such tools are united by common design principles and theories of online human behavior, or whether they reflect disparate assumptions. This study fills this gap in scholarly knowledge about credibility tools through an iterative mixed-methods UX research and design approach. Through interviews with both designers and users of credibility tools, usability studies, and survey experiments, we seek to formulate a common UX framework for the standardized design and development of digital news credibility tools. Bearing in mind the specific challenges that socially networked digital media and online attention dynamics pose for UX design, we name this the “UX of Credible Content” framework. From this framework, we plan to design a prototype of a credibility tool that synthesizes best UX practices and insights.

Introduction

As public awareness of online mis- and disinformation has increased in recent years, a number of tools and technological interventions have been developed to help users better identify manipulative content and publishers, and to evaluate the credibility of online information at large. Most commonly, such efforts manifest as third-party web browser plug-ins that allow users to quickly surface information and context about a particular publisher or piece of content (Figure 2).

While a considerable body of research exists on different sociotechnical methods for evaluating the credibility of online content and countering misinformation, significantly less literature has explored the human factors and user experience (UX) design considerations for credibility tools. How can credibility tools be designed to augment and support human judgment and media literacy, while realistically

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

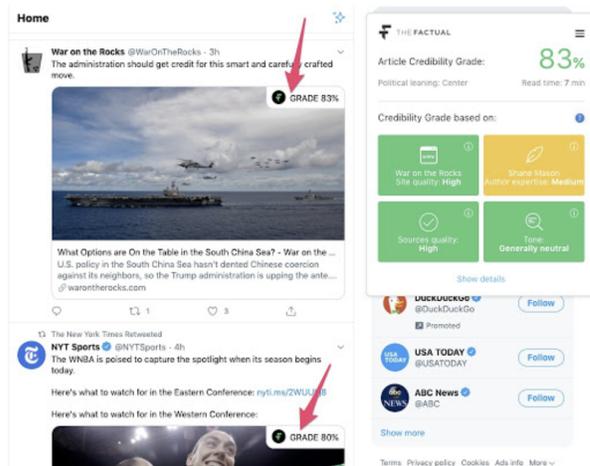


Figure 1: The user interface of The Factual, an example of a browser extension credibility tool (The Factual, 2020)

acknowledging the conditions of the online attention economy? Taking into account cognitive biases, heuristics, and human irrationality—all of which influence how users process and evaluate online information—how can credibility tools be designed to best support human reasoning, judgment, and decision-making when encountering online content? How do the back-end methods of credibility evaluation influence the front-end interface design possibilities and considerations? For example, how should probabilistic or uncertain algorithmic outputs be communicated to users in a way that builds healthy trust but discourages over-trust? What design features are necessary to ensure the fairness, accountability, and transparency of credibility tools from the user’s perspective?

These questions are critical, given that current research suggests that the efficacy of current social media labeling interventions—the platforms’ attempt to introduce credibility signals—are low and may even occasionally deepen distrust of the platforms (Emily Saltz and Wardle 2020). Moreover, the use of algorithms and AI introduce additional complications to establishing user trust, given the unique challenges to fairness, accountability, and transparency that the use of machine learning and algorithmic intelligence poses.

The emerging field of explainable AI (XAI), however, responds to these challenges, through rigorous inquiry into both the technical and psychological elements that AI explanations must consider (Brian Lim and Ashraf Abdul 2019; Morris et al. 2012).

A key question for both automated and manual credibility tools, however, is how context might affect the conditions for a satisfactory explanation of the credibility evaluation, especially given the role emotional and cognitive biases play in media consumption. Without a standardized credibility scale or framework, credibility tools must rely on third-party consensus or persuade users with their own credibility evaluation methodology. In either case, credibility tools face significant challenges to gaining the user's trust, especially when the judgment conflicts with the users' expectations or beliefs. In such cases of expectation violation, transparent and easily interpreted explanations that invoke fair and unbiased reasoning will be critical to persuasion. Credibility tools may further demonstrate their accountability in this situation by establishing a feedback mechanism through which users may contest judgments and provide their reasoning. Many online news credibility tools incorporate some degree of automation and thus face an additional set of challenges to user trust: technical literacy and algorithmic opacity. Such tools need to balance competing needs to make their credibility evaluations trustworthy to users and may require a unique framework that builds off of XAI principles but incorporates media psychology research.

While there have been efforts to comprehensively catalog the emerging tools aimed at fighting online misinformation, there has not yet been any systematic evaluation or investigation of such tools. Recent research has suggested that news source credibility overlaps yet varies significantly "in the eyes of different assessors." (Mensio and Alani 2019, p.2) While such diversity of evaluation may be useful to "reduce bias and encourage further debate," it is an open question how cases of disagreement among assessors should be handled and responded to; the lack of a standardized credibility scale further compounds this issue.

In addition to the issue of the diversity of credibility assessments, there is the issue of the diversity of design choices. It is unclear which credibility tools are actually effective and useful to their users, and moreover, if these tools share common design frameworks and goals, or are based on unique or disparate assumptions made by individual organizations and designers. We fill this knowledge gap through mixed-methods study of existing credibility tools, interviews with both tool designers and users, and through design studies of experimental prototypes. Our findings through these studies will manifest the UX of Credible Content design framework which will aid designers and developers who attempt to introduce tools and technologies assessing the credibility of digital news content.

Related Literature

Several recent psychology experiments have been conducted that measure the empirical effects of content warnings, fact-check labels, and credibility indicators on social media content sharing (Clayton et al. 2019; Ecker et al. 2019;

Yaqub et al. 2020). Yet there are no universal standards that guide the design and language of their tested annotations. In a sense, therefore, these studies are measuring different things.

These experiments also fail to account for the roles of user experience (UX) and interface design, which are known to have strong effects on trust and perception. Multiple studies find that users put more emphasis on factors like visual design elements and navigability of websites than the actual information presented when assessing credibility (Metzger and Flanagin 2013; Fogg et al. 2003). Thus, labeling interventions that do not make use of modern web design and aesthetic standards are unlikely to demonstrate their full potential efficacy.

Despite the limitations and design disparities, the results from different labeling intervention trials still provide useful insights that can be triangulated. Ecker et al. find modest evidence that short-format (140 character) refutational fact-checks are more effective than simple 'false-tag' refutations and that both reduce belief in false claims (Ecker et al. 2019). Clayton et al. find modest support for simple tags ("Disputed" and "Rated false") but show that general warnings about misleading information on social media can backfire by reducing belief in true headlines (Clayton et al. 2019). This finding is consistent with prior research that theorizes an "Implied Truth Effect," whereby the presence of warning labels causes unlabeled stories to be seen as more accurate (Pennycook et al. 2020). It is unclear, however, why this theory would not apply equally to the simple tag refutations.

Yaqub et al. compare four 1-sentence disputing labels, evaluating how different judgment attributions (artificial intelligence, public consensus, news media, or fact-checkers) affect the label's efficacy (Yaqub et al. 2020). They find that "the Fact-Checkers indicator was the most successful in persuading people to avoid sharing Non-true headlines" with a 43 percent reduction in sharing intent, whereas the artificial intelligence (AI) indicator was the least effective, with only a 22 percent decrease. This evidence presents significant implications for both existing and future platform mediation interventions, which typically rely on the combination of human and algorithmic labor to successfully implement. Here, we borrow Caplan's definition whereby content moderation is the process of "identifying and removing content that is demonstrably false and/or harmful," whereas content mediation is the process of "identifying what information is credible, when, and how to communicate these changes" (Caplan 2020, p.1). While human fact-checkers ultimately decide what is credible, AI is often used to filter and label content instantaneously at scale. Both platform moderation and mediation often invoke machine learning or some other form of automation.

Currently, platforms only deploy this type of automated detection when technically necessary to the efforts of content review teams, however, it could also power user-facing tools that augment public judgment. Arguably, users deserve better signals of information quality for all content, not just claims determined relevant by self-interested private companies. If mixed-initiative human-algorithmic labor can detect and evaluate information credibility at scale, how might it

best support human media literacy? In light of the myriad deficits in public and institutional trust, how can such a sociotechnical system be made trustworthy?

The biggest limitation to existing labeling interventions is their oversimplification of information credibility. A narrow focus on refutation implicitly endorses a binary view of factuality, which is unable to respond to more complex and ambiguous cases. Not all misinformation is equally wrong: the number of falsehoods, the scale of inaccuracy, and the relative importance of erroneous claims all factor into how “wrong” an article is considered to be. It is possible to deceive with verifiably accurate information: the intentional exclusion of information, evidence, and context is as important to credibility as information cited. Both factuality and credibility exist on a spectrum.

Credibility, at its core, is a measure of trustworthiness, and rooted in the truster’s holistic perception of the source, the social context, the communication medium, and the information itself. While most misinformation interventions and warning labels have focused on validity and verification – the factuality or authenticity of specific claims or media within content – credibility incorporates additional elements that influence the believability of content, including the source, bias, number and quality of citations, logical soundness, and adherence to journalistic guidelines (Zhang et al. 2018). Modern theories of credibility incorporate elements specific to digital media as well: for example, Sundar advocates for the MAIN model of web credibility, whereby a webpage is judged by cues corresponding to its modality, agency, interactivity, and navigability (Sundar 2007). Ultimately, a credibility framework is capable of responding to the entire content spectrum and is resilient to gamification by adversaries when several or more credibility signals are incorporated. While further research is needed to establish which credibility signals are most significant and reliable, the credibility tools we test focus on basic and widely agreed upon credibility indicators: source, author, number and quality of citations, publication history, and history of fact-checks.

Cognitive psychology research and theories of information processing provide reason to believe that holistic credibility evaluations may be more persuasive than simple fact-checks. As Kahneman’s account of the dual-process theory of reasoning posits, our relationship to information is primarily emotional and only secondarily rational (Kahneman 2013). We tend to base our initial judgments on emotions and cognitive heuristics (system 1 thinking) and are prone to justifying them with post-hoc rationalization (system 2). Indeed, many media literacy interventions simply aim to nudge people to engage in more deliberative thinking to limit the influence of system 1’s intuitive impressions on conclusions (Bence Bago and Pennycook 2020). After people have already formed a strong emotional attachment to a source or idea, fact-check refutations are not only prone to failure but may even deepen distrust or negative sentiments of the intervening authority (Emily Saltz and Wardle 2020). Acknowledging subjective reality, “modern accounts of credibility define it as the believability of a source, and it rests largely on perceptions of the trustworthiness and ex-

pertise of the information source as interpreted by the information receiver.” (Metzger and Flanagin 2013, p.211)

Credibility evaluations and fact-checks are unlikely to persuade individuals emotionally attached to misinformation. Clear signals of credibility deficits, however, might prevent the formation of new attachments to misleading sources. Such signals can incorporate both appeals to system 1 (source, context, tone, author expertise) and system 2 (factuality, logical soundness). Nonetheless, these signals must be analyzed, distilled, and made salient; individuals generally lack sufficient motivation to process them on their own. Due to bounded rationality and limited cognitive capacity, people typically maximize their reward to effort ratio when processing information by leveraging the minimum cognitive resources necessary to provide an outcome sufficient to their needs (Geeng, Yee, and Roesner 2020) When factoring into account the structural properties of networked digital media that further burden individual cognition, it becomes clear why “using peripheral cues” to judge credibility is “the rule of web use, not the exception.” (Fogg et al. 2003, p.15) Credibility signals, therefore, must be rendered as peripheral cues to users.

Social media companies may currently lack incentive, however, to implement such credibility signals. Media organizations, after all, are clients of platforms; they pay platforms to advertise their content. Both parties are rewarded when content attracts audience engagement. Consequently, implementing signals that may deter audiences poses financial risks: potential loss of revenue and or client partnerships. Further research is needed, however, to understand the effects of credibility signals on overall platform and content engagement, and how financial risk can be minimized. Moreover, platforms are rightly held partly responsible for the soundness of judgments they superimpose onto content; they have reason to limit their use of credibility labels – especially those generated by ML, which by nature of probabilistic inference, cannot provide perfect certainty nor causal evidence. Platforms thus benefit from more restrained approaches to both mediation and automation, deploying labels only when deemed urgently necessary by context.

Third-party organizations, however, have begun to explore more diverse and holistic interventions by building their own credibility tools, commonly in the form of web browser extensions. Typically, such extensions “surface information about a news article to benefit readers, much like how nutrition labels for food and browser security labels for webpages provide context in the moment.” (Zhang et al. 2018, p.604) This initiative relies on forms of human judgment, but several tools have successfully automated the process with AI. While third-party tools differ from platform-based interventions in their opt-in nature, the UX design principles for credibility tools apply equally platform-based interventions; when considering only the interfaces, interactions, and affordances for users, credibility tools and platform-based interventions are in essence the same: they alter the webpage markup to alter users’ perceptions of content, ideally empowering the user to make a more informed choice. In this regard, third-party credibility tools – with their unique affordances and lack of constraints

that platforms have – may serve as design experiments for tools and interfaces that platforms or browsers may eventually adopt, or be built with. Third-party approaches also benefit from decentralization, allowing a greater diversity of intervention approaches and greater user choice, diffusing the power of the platform as an intervening authority. However, a drawback to this fragmentation is the lack of standardization [of credibility evaluation and tool design] leading to greater confusion in the face of choice overload or lack of consensus. Moreover, it may be more difficult to trace responsibility for credibility evaluations and indicators without standardization.

The development of credibility tools tools, while encouraging, has effects on users’ perceptions and behavior and holds implications for UX design studies that are not well understood and largely absent in scholarly literature. As stated earlier, while a diversity of technical and social fact-checking and credibility evaluation methods have been studied, the front-end design of tools that implement these methods is just as important; credibility tool design must consider this specific context of human-computer interaction to be made useful in practice (Nguyen et al. 2018).

Tool Selection Process

The process of software selection began with the attempt to conduct a complete search for all of the publicly available fact-checking and credibility evaluation tools that could be found online. Specifically, tools were found through exhaustive searches of the Google Chrome and Firefox web browser extension stores, as well through two databases of disinformation web tools compiled by RAND corporation and the Credibility Coalition’s UX Working Group. Given the aforementioned principle of cognitive miserliness and the influence of intuitions and automatic reflexes on decision-making, we narrowed our consideration to tools that automatically inject labeling and or other immediate visual signals via alteration of the webpage markup [as opposed to web apps or search portals], which are all web browser extensions by nature of technical affordance. The appropriate tools were then each tested by the first author systematically to determine whether they passed a minimum level of accuracy, reliability, and usability. Tools that malfunctioned or were consistently inaccurate or unintelligible were discarded. After further review, 4 tools were selected for usability study that represent conceptually distinct approaches to credibility evaluation and intervention method: a source-based system (NewsGuard), a tone/text-based system (FakerFact), a hybrid source/text system (The Factual), and a media-bias focused system (Media Bias Fact Check Extension). Moreover, these tools differ in their degree of automation and use of machine learning ML); for example, FakerFact relies entirely on natural language processing (NLP) to classify content, whereas The Factual uses ML algorithms trained on credibility annotations to generate credibility scores, and NewsGuard relies exclusively on human expert evaluation. Each of these tools offer detailed explanations for their credibility judgments and provide differing levels of transparency and contextual information to users, however, they differ in their UX of con-

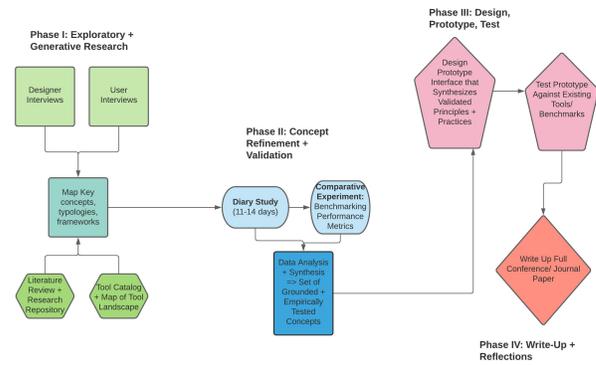


Figure 2: A flowchart of our research plan, comprised of a generative research phase, an evaluative research phase, and a final phase of iterative design and user testing.

veying this information: for example, the Media Bias Fact Check entirely removes the image and markup for ‘problematic’ sources, whereas NewsGuard only inserts a small red ‘x’ in the preview markup for publishers deemed problematic. Given their significant differences in all aspects of UX design, we believe this tool sample will result in illustrative differences that can be abstracted to broader design principles.

Methodology

Our study follows a design similar to that of an exploratory, sequential mixed-methods approach; an exploratory, generative phase of broad, qualitative research and data collection is followed by an evaluative period of research that seeks to refine, test, and validate insight gathered from the exploratory phase. The final phase of research synthesizes the findings and insight from the generative and evaluative research towards potential design frameworks and solutions. (Figure 2)

The first, generative phase of our research is based on semi-structured, hour-long interviews to gather deep, qualitative insights about the nature of the problem of online misinformation and credibility evaluation, the specific challenges from this problem that designers face, and the needs this problem creates for users that have not been met by designers yet. To achieve a fuller understanding of the design of media credibility tools, we will gather data from two diametric perspectives: that of UX designers, and that of potential users. For a broader perspective of the principles behind designing for a UX of credibility and against misinformation, we will interview both UX designers of third-party tools as well as platform-based interventions. Designers’ assumptions and intentions shed light on why and how media credibility tools were designed; in contrast, users’ perceptions and emotions help us understand how well designers’ intentions are received, and how valid the designers’ assumptions are. While conducting interviews in this first stage of research, we will also be compiling as much knowledge as we can on past and current approaches to the design

of credibility tools, through literature searches and searching the web for existing tools. We will integrate this knowledge with the qualitative analysis of our interview data to contextualize the results. By studying both UX designers' intentions and users' perceptions of media credibility tools through qualitative research, we aim to develop a more holistic and standardized design framework and set of principles for online credibility tools.

The second phase of our research involves refining and evaluating our framework through user-testing of existing tools and a diary study over a longer period in a natural, real-world context. In online survey experiments, we will comparatively user-test our tool sample, which we believe to represent distinct intervention and UX approaches, to benchmark for both usability metrics and their actual effects on user perception of media. To achieve this, we will have present survey participants with an artificial social media 'newsfeed' of headlines with markup images and text. First, participants will rate the believability and trustworthiness of each headline in the newsfeed unlabeled by credibility tools. Then, the participants will be randomized into one of four treatment groups, each of which will have the same newsfeed, but annotated with labels by one of the four credibility tools, and will then be asked to re-rate the believability and trustworthiness of the content with the credibility tool's labels present. After the experiment, participants will be asked overall how much they trusted the evaluations offered by the credibility tools, and will be offered a free response opportunity to indicate what they wish had been different or would need to see to trust the tool's judgment. Analyzing the data, we will compare how the tools are consciously trusted by users, as well as how much of an influence they had on user perception of the newsfeed content.

Our diary study will complement our survey experiment by offering greater ecological validity [through a context of real-world tool use] and a longer time-span of study than a survey experiment can afford. Our key questions motivating the diary study are: 1) do users still engage with the credibility tools when they are not consciously asked to do so, and have no immediate artificial incentive to do so? and 2) do users still consciously engage with credibility tools when they are no longer novel, or do they adapt to their presence and overlook them over time? To answer this, we will ask our users to document their daily engagement with the installed credibility tool as honestly as possible, including days zero interaction occurs, by having them journal nightly over a 10-day period via the dscout user research platform. We will also ask our participants to document or screenshot any interactions that are particularly salient or problematic, or represent a pivotal moment in their relationship to the tool. This will allow us to capture a wider range of potential usability and design issues than we are able to in an artificial experiment or from our own internal testing. As many of the credibility tools are limited to annotating publishers that the tool has collected sufficient data on, a particular question is whether existing credibility tools are unable to benefit certain user groups with an affinity for user-generated or non-mainstream content. Indeed, given the rising primacy of memes and user-generated content, and the resulting shift

from text to multimedia, we suspect that existing credibility tools may be limited in their range and require a redesign that doesn't overly emphasize classic, text-based news articles.

For each step of our research plan targeting public end-users, we aim to recruit a sociodemographically diverse sample of 18-70 year-old, English-speaking, social media users, roughly representative of the United States population. We will ask for self-reported political-leaning and take into account the particular effect of this variable in all of our analyses.

The third and final phase of our research plan is to design a working prototype of a credibility tool based on our research insights and UX of credible content framework, one that incorporates the best practices and principles we have observed through our mixed-methods inquiry. After arriving at a functional prototype through iterative design and testing of initial sketches and wireframes, we will conduct mixed-methods user tests and experiments with the prototype to gather further empirical feedback and validate our findings. Ultimately, the goal is to productively synthesize our findings from a holistic investigation of media credibility tools to create empirically validated and actionable design guidelines and practices. We hope to offer both the public and the UX design community a working prototype as proof of concept and a point of inspiration for further design studies.

Conclusion

The issues underlying the contemporary 'post-truth era' are larger than just the presence of junk news. The constituent problems are complex and multi-sited, involving the interaction of sociological, technological, structural, and systemic factors, which each exert influence at both the individual and collective levels. There is, therefore, no single solution nor single organization capable of 'fixing' the post-truth, social media era. Confronting the online causes of the contemporary post-truth era requires collaboration between social media companies, governments, policymakers, academics, and third-party technologists. Fighting misinformation necessitates interventions aimed at its publishers and promoters to regulate the policy, platform design, and business practices of social media companies, and interventions targeting the public, e.g., individual social media users. Credibility tools, however, may play a significant role in offering crucial context and background information to users that may be otherwise baited by deceptive sources. While the existing credibility tools studied here may have little impact because of their voluntary effort required to find and install the tools, in the future, if better integrated into users' experience, they could become an important part of the software ecosystem as 'middleware'. Moreover, credibility tools offer an experimental means of studying UX design principles that may ultimately guide the future design of ethical browsers or social media platforms, allowing us to ascertain what information is most important to be presented to consumers of networked digital media, and how it should be presented.

Acknowledgments

This work is grateful for the support of the Credibility Coalition and Hacks/Hackers. It is also thankful to the University of Washington for its ethical review and approval.

References

- Bence Bago, D. G. R., and Pennycook, G. 2020. Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *Journal of Experimental Psychology* 149(9):1608–1613.
- Brian Lim, Q. Y., and Ashraf Abdul, D. W. 2019. Why these explanations? selecting intelligibility types for explanation goals. Joint Proceedings of the ACM IUI 2019 Workshops.
- Caplan, R. 2020. Covid-19 misinformation is a crisis of content mediation.
- Clayton, K.; Blair, S.; Busam, J. A.; Forstner, S.; Gance, J.; Green, G.; Kawata, A.; Kovvuri, A.; Martin, J.; Morgan, E.; and et al. 2019. Real solutions for fake news? measuring the effectiveness of general warnings and fact-check tags in reducing belief in false stories on social media. *Political Behavior* 42(4):1073–1095.
- Ecker, U. K.; O’Reilly, Z.; Reid, J. S.; and Chang, E. P. 2019. The effectiveness of short-format refutational fact-checks. *British Journal of Psychology* 111(1):36–54.
- Emily Saltz, C. L., and Wardle, C. 2020. Encounters with visual misinformation and labels across platforms: An interview and diary study to inform ecosystem approaches to misinformation interventions. arXiv.org. <https://arxiv.org/abs/2011.12758>.
- Fogg, B. J.; Soohoo, C.; Danielson, D. R.; Marable, L.; Stanford, J.; and Tauber, E. R. 2003. How do users evaluate the credibility of web sites? a study with over 2,500 participants. DUX ’03, 1–15. New York, NY, USA: Association for Computing Machinery.
- Geeng, C.; Yee, S.; and Roesner, F. 2020. Fake news on facebook and twitter: Investigating how people (don’t) investigate. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*.
- Kahneman, D. 2013. *Thinking, Fast and Slow*. Farrar, Straus, and Giroux.
- Mensio, M., and Alani, H. 2019. News source credibility in the eyes of different assessors. Conference for Truth and Trust Online.
- Metzger, M. J., and Flanagin, A. J. 2013. Credibility and trust of information in online environments: The use of cognitive heuristics. *Journal of Pragmatics* 59:210–220. Biases and constraints in communication: Argumentation, persuasion and manipulation.
- Morris, M. R.; Counts, S.; Roseway, A.; Hoff, A.; and Schwarz, J. 2012. Tweeting is believing? understanding microblog credibility perceptions. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, CSCW ’12*, 441–450. New York, NY, USA: Association for Computing Machinery.
- Nguyen, A. T.; Kharosekar, A.; Krishnan, S.; Krishnan, S.; Tate, E.; Wallace, B. C.; and Lease, M. 2018. Believe it or not: Designing a human-ai partnership for mixed-initiative fact-checking. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology, UIST ’18*, 189–199. New York, NY, USA: Association for Computing Machinery.
- Pennycook, G.; McPhetres, J.; Zhang, Y.; Lu, J. G.; and Rand, D. G. 2020. Fighting covid-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. *Psychological Science* 31(7):770–780. PMID: 32603243.
- Sundar, S. 2007. The main model : A heuristic approach to understanding technology effects on credibility.
- Yaqub, W.; Kakhidze, O.; Brockman, M. L.; Memon, N.; and Patil, S. 2020. Effects of credibility indicators on social media news sharing intent. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI ’20*, 1–14. New York, NY, USA: Association for Computing Machinery.
- Zhang, A.; Ranganathan, A.; Metz, S. E.; Appling, S.; Sehat, C. M.; Gilmore, N.; Adams, N. B.; Vincent, E.; Lee, J. ; Robbins, M.; Bice, E.; Hawke, S.; Karger, D.; and Mina, A. X. 2018. A structured response to misinformation: Defining and annotating credibility indicators in news articles. International World Wide Web Conference (WebConf) Companion.