Federal Register Notice 87 FR 15274, https://www.federalregister.gov/documents/2022/03/17/2022-05683/request-for-information-on-federal-priorities-for-information-integrity-research-and-development, May 15, 2022

Request for Information on Federal Priorities for Information Integrity Research and Development

SIFT, LLC

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views or opinions of the U.S. Government. We bear no responsibility for the accuracy, legality, or content of the responses and external links included in this document.

RESPONSE TO: REQUEST FOR INFORMATION ON FEDERAL PRIORITIES FOR INFORMATION INTEGRITY RESEARCH AND DEVELOPMENT

SIFT, LLC

Misinformation and disinformation are longstanding problems that have recently grown in complexity due to enabling technologies. The lines of transmission are more numerous, rapid, and dynamic than ever. We propose technology and research questions aimed at empowering the characterization of information reliability, the production of more reliable information, as well as permeating the public square with new literacy, expectations, and standards related to information integrity. The central conceptual pillar of our approach is information provenance – structured data about how the information came to be. As one concrete application of provenance-based analytics, SIFT has developed Project7, an experimental, collaborative human-machine information analysis workspace. Project7 demonstrates that provenance graph representations and interactive provenance-based displays can help assess the origins, bias, confidence, and integrity of information. Its development has been informed by standards, directives, and metrics for integrity and rigor originating from the intelligence community. We propose to utilize information analysis workspaces such as Project7 as platforms for investigating information integrity research questions spanning intelligence analysis, journalism, and other content creation domains. Research in this area has implications for characterizing and increasing information integrity and improving consumers' resilience to misinformation and disinformation.

Introduction

Misinformation and disinformation are problems older than the printing press, but in the digital age they have swelled into massive challenges that hinder our ability to establish a commonly held set of basic facts with which to collectively debate and reason. The explosion of information pathways has rendered our previous defense mechanisms, which centered on fostering analytical rigor and integrity in the institutions with the most reach, insufficient - yet more important than ever. In this response we propose promising leads and timely questions that aim to raise the bar of informational integrity to a new high for all information producers, brokers, and consumers, by empowering them with tools and standards for analytical rigor.

Two critical efforts for mitigating the effects of misinformation and disinformation are (1) advancing the trustworthiness of information, and (2) enhancing the *informational immune systems* (i.e., their ability to discern and resist mis/disinformation) of individuals who consume information, drive commercial incentives for information producers, disseminate information through social networks, and increasingly participate in the production of journalism. These two pursuits are closely related, and we briefly discuss some connections between them before describing technological advances.

Trustworthiness & Perceived Credibility

For information consumers, a key asset is the availability of sources that are both *perceived* as trustworthy and are also *worthy* of that trust. Guillory and Geraci (2013) found that incorrect initial inferences, while powerful and persistent, can be corrected if the correction comes from a trustworthy source. In that study, trustworthiness (i.e, the perceived willingness to be accurate) was found to be more persuasive than expertise (i.e., the perceived ability to be accurate), which agrees with previous studies (McGinnies & Ward 1980; Lui & Standing 1989). Perceived credibility – closely related to but not synonymous with trustworthiness – has also been found to be associated with a higher likelihood of future engagement with a news source (Peifer & Meisinger 2021), which suggests that establishing credibility when disseminating information can have lasting benefits on impact that stretch beyond that piece of information alone. This evidence converges with the positive link between trust in a news source and loyalty to it, found by Nelson & Kim (2021). All else being equal, more trusted information brokers have greater potential to become resources for vetting, contextualizing and ultimately neutralizing misleading information. Broad-based trust is powerful, but rare because it is so arduous to build and so susceptible to self-destruction.

Improving Credibility, Reliability, and Transparency of Information Producers

One intuitive way to improve the credibility of information producers is to help them be more reliable (i.e., more accurate more often) by creating tools that enable them to conduct analyses with rigor (i.e., thorough, clear, and able to be validated and critiqued for confidence) (Zelik et al., 2010). Here we see a valuable role for software assistants that help track and publish *provenance* – data about the information's synthesis and the sources from which it came – and assist with critical-thinking tasks using that provenance. We illustrate this in later sections.

It would be ideal for these tools to also allow information producers (e.g., reporters, intelligence analysts, and their software analytics) to be able to "show their work" without compromising sensitive sources and methods, to facilitate better collaboration and transparency with consumers. For two decades, since the burgeoning of online news, blogging and the adage "transparency is the new objectivity," much research has gone into investigating the effect of transparency on credibility and trust in journalism. Some results have confirmed that there exists a positive relationship (Curry & Stroud 2019; Peifer & Meisinger 2021), but others suggest a more nuanced relationship. Karlsson and Clerwall (2018) found hyperlinks to be a form of transparency that was met with particularly positive response from readers because "hyperlinks make it possible to track down original sources and documents," among other reasons, while providing "negative user commentary" as transparency was counterproductive. Tandoc and Thomas (2017) found a negative effect of transparency on credibility when the form of transparency was the disclosure of biographical information about the author. Karlsson (2020) found that "participatory transparency" – the involvement of users in various stages of the reporting process, e.g. by sending pictures of events –increased source trustworthiness with the group of readers that had the most skeptical attitudes towards news media (low educated males), while it decreased trustworthiness with the least skeptical group (highly educated females). Transparency's real-world implementation is also a complicating matter. Despite transparency's efficacious origins in the blogging world (Lasika 2004), Koliska and Chadha (2016) later studied newsrooms and found that the practice of transparency in corporate journalism had become

largely performative, having been mandated and implemented without the direct involvement of journalists. A common consensus of the research into transparency's effect on credibility seems to be that "transparency" encompasses too many elements for sweeping conclusions to be made, and that different audiences favor different forms of transparency.

The expectations and effectiveness of different transparency and disclosure strategies also may be highly dependent on the underlying analysis and reporting processes. This is consistent with the finding of Diakopoulos and Koliska (2017) that information consumers desire certain types of transparency from algorithmic (i.e., machine-generated) reporting, including information about the data, the model, inference methods, and the availability of a public-facing interface into these aspects.

The previous work described above suggests that trustworthiness of information – and consumers' ability to characterize it – can be enhanced by tools that support (1) more rigorous, collaborative analysis – a form of quality assurance – and (2) customizable transparency about the genesis of the information, including machine reasoning. These are some of the primary motivations behind SIFT's ongoing R&D on provenance and the Project7 human-machine analysis workspace. Below we discuss these technologies, their relevance to intelligence analysis and journalism, and relevant standards and metrics for analytical rigor. We outline some connections between these tools and ideas for enhancing the public's immunity to manipulated information, identifying relevant research questions.

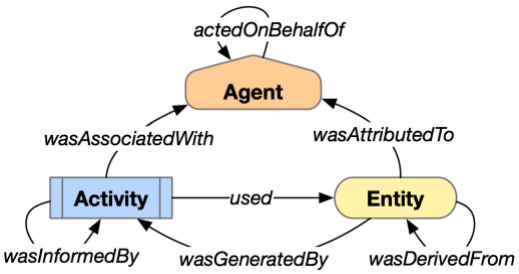

Figure 2: The PROV data model.

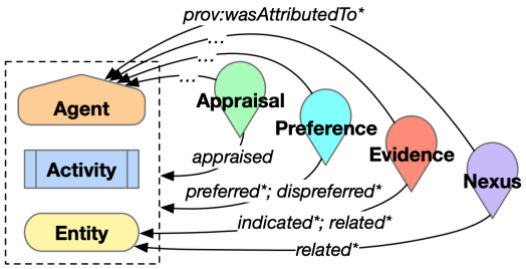

Figure 1: The DIVE ontology.

Provenance

The provenance of a piece of information refers to data about its origins. Formal provenance ontologies consist of the types of – and types of relationships between – elements that can play a role in these origins.

SIFT's provenance research products are built around the W3C-recommended PROV data model (Figure 1), which contains three types of elements: entities, activities, and agents. *Entities* are fixed real or hypothetical things, such as records, assertions, databases, etc., which can be input to or outputs of activities. *Activities* are processes that occur over a period of time, such as inference actions or other procedures performed by human or machine. *Agents* are those actors that perform activities, whether they be humans, organizations, web services or machine learning modules.

The predicates in the PROV data model relate particular types of elements to other types of elements. PROV can represent that an entity **was derived from** another entity, **was attributed**

to an agent, or **was generated by** an activity. It can represent that an activity **used** an entity, **was informed by** another activity, or **was associated with** an agent. Finally, it can represent that an agent **acted on behalf of** another agent. All of these relationships help represent directed information dependency, e.g., an entity that **was generated by** an activity can be seen as immediately *downstream* from that activity, and that activity would in turn be immediately downstream of any entities that it **used**. In this way, the PROV data model can be used to represent the full inferential provenance for a piece of data (e.g., an assertion) as a dependency graph, called a provenance graph.

Provenance alone helps express the structure of an analysis, but it does not express the information integrity considerations of information diversity, confidence, alternatives, assumptions, gaps, conflicts, likelihood, or biases. To help express these considerations, SIFT and BBN extended the PROV data model with the DIVE ontology (Friedman et al. 2020; Figure 2). DIVE adds four classes of provenance elements, all of which represent judgments that are attributed to agents (**wasAttributedTo**). An *Appraisal* is a confidence judgment about any other (**appraised**) element, with attributes for confidence, likelihood, bias, and reliability. A *Preference* is a judgment about the relative quality between one (**preferred)** element and another (**dispreferred**) element, all else being equal. *Evidence* is a judgment about the diagnosticity of one (**related**) entity on another (**indicated**) entity, e.g., evidence toward a hypothesis. A *Nexus* is a judgment about the mutual coherence or conflict within a set of (**related**) entities (i.e., that the entities in set have high or low joint likelihood). These extensions capture local quality judgments within the provenance graph that can support global information integrity assessments. Follow-on research should extend DIVE to cover a fuller range of human explanation and argumentation, in such a way that fits into a dependency graph.

The provenance graph can serve as a substrate for critical thinking about sensitivity, confidence, information necessity and sufficiency, and impact, using ATMS-inspired algorithms (Forbus & de Kleer, 1993; Friedman et al. 2021). For example, new assessments about the reliability of sources upstream can potentially justify shifts in judgments about the quality of resulting inferences downstream. Similarly, semantic information tags such as INTs (*OSINT***,** *HUMINT***,** *IMINT***,** etc.) source types (e.g., *Online News*, *Social Media*), operation types (e.g., *NLP*, *Named Entity Recognition***,** *Machine Learning*), and operating assumptions (e.g., inferring a vessel's location via a transponder signal implies the transponder is *on* the vessel) may be added to individual nodes, and algorithms can propagate these downstream. These propagation algorithms allow the user to (temporarily) remove a source to assess its downstream impact on their conclusions, propagate confidence downstream from sources to estimate the confidence of conclusions, and assess the diversity, gaps, and assumptions in downstream conclusions.

Provenance Graphs in Project7

Project7 is an experimental human-machine information analysis workspace supported by multiple efforts for high-integrity information analysis (Friedman et al. 2021). Project7 allows multiple users to collaborate on the generation and validation of competing hypotheses, with an interactive view of the provenance graph playing a central role.

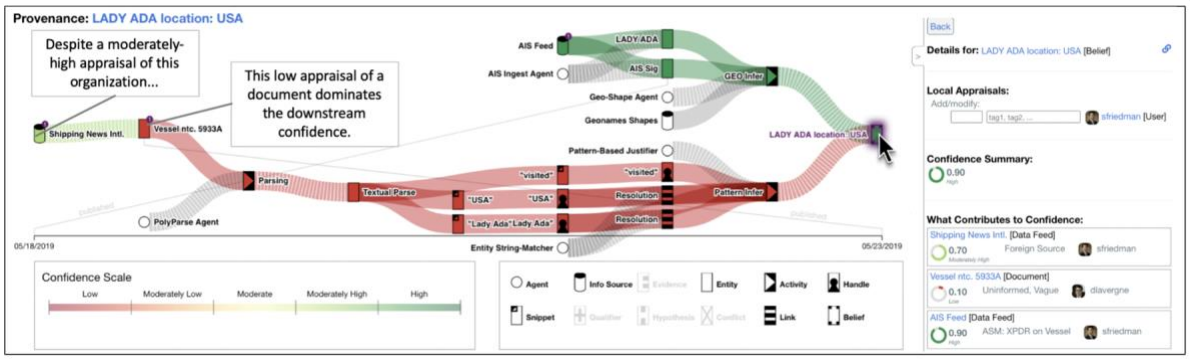

Figure 3: A provenance graph in Project7 for the inference that the ship LADY ADA was located in the USA.

Figure 3 shows an example provenance display for the hypothesis that the fictional vessel *Lady Ada* is located in *USA.* The hypothesis (Figure 3, right) was generated by two alternative inferential paths through the provenance, one with a high confidence and the other with low confidence, resulting from user appraisals upstream. Project7 allows users to explore alternative, experimental approaches to propagating quality judgments through the graph. In Figure 3, confidence propagates rightward through junctions as if they are an AND (min input value) or an OR (max input value) depending on whether the incoming relations are semantically necessary vs. sufficient. The interface allows users to conduct sensitivity analysis by excluding elements by their identity or properties, including users. In this example, excluding the user *sfriedman* would invalidate their appraisals for *Shipping News Intl.* and the *AIS Feed*, and the high confidence would cease to flow through the upper path, ultimately dropping the confidence in the hypothesis to low (via only the bottom path). This helps users quickly answer *"what if?"* questions and investigate the global impact of particular elements in the analysis.

Using provenance as an interactive substrate facilitates user-driven exploration of the sensitivity of inferences to hypothetical changes in their related assumptions, processes, data sources, and key contributors. Expanding this approach and applying it more broadly has the potential to revolutionize how content is collaboratively generated and disseminated with a suitable level of transparency. An open research question is how the machine can proactively contribute to this exploration, acting as a critical-thinking assistant and radically improving rigor by mitigating cognitive biases with low cost, high coverage, and high throughput.

Intelligence and Journalism

Intelligence analysis and journalism share many of the same concerns and best practices. Both enterprises collect information, and, where possible, synthesize meaningful inferences for decision makers in the information ecosystem. Both have access to specialized sources and collection methods with diverse capabilities and technical risks. Both have reasons to "show their work" (assumptions, sourcing, argumentation) as much as possible while also protecting sources and methods. Both practices can be collaborative, which adds communication risks to the reporting process, and creates an imperative to keep fulsome and explicit records.

A significant difference is that journalists' output, when ready, is typically dispersed broadly into the public information sphere, whereas intelligence analysts' outputs have a more

variegated audience and distribution methods governed by security controls. The provenance graph is itself a useful data structure for propagating these controls from sources and methods to assessments. In this dimension, we see the practice of journalism as subsumed by the practice of intelligence analysis, and thus by starting as an intelligence analysis workspace, ideas from Project7 are well positioned to generalize to journalism. In both domains, information integrity tools can help maintain standards for rigor in an environment where human operators are susceptible to intrinsic cognitive biases, incentive structures, and attentional limitations.

We posit that a long-term goal for building resiliency to misinformation and disinformation should be to enhance the public's literacy, expectations and tools for provenance and rigor in the information they consume. This shift can be driven in part by information producers, with the help of our research and development.

Supporting Rigorous Analysis

Provenance-based analytics can advance the trustworthiness of information by supporting aspects of rigor for analysts in the intelligence community and journalism. Here we draw attention to noteworthy properties of analytic rigor that motivate SIFT's provenance-guided analytics, and to others where analytics like Project7 are not yet sufficient.

The Office of the Director of National Intelligence (2021) has published intelligence community directives (ICDs). ICD 203 contains nine tradecraft standards intended as guides for achieving analytic rigor and quality. Similarly, Zelik et al.'s (2010) eight attributes of analytical rigor also provide guidance for high-quality information analysis. We review these standards and metrics of rigor to illustrate how provenance addresses integrity and rigor in this domain.

ICD 203's standard two advises that intelligence "Properly expresses and explains uncertainties associated with major analytic judgments." As discussed above, provenance-based analytics allows users to appraise elements in the graph with a confidence scale that has quantitative and qualitative aspects, and these confidences propagate through the graph using a set of alternative propagation schemes.

Standard three advises that an analytic product "Properly distinguishes between underlying intelligence information and analysts' assumptions and judgments." As shown in Figure 3, Project7 uses iconography to display canonical graph nodes (e.g., beliefs) and semantic tags, including underlying assumptions, to express source diversity and facilitate visual filtering.

Standard four, and Zelik et al.'s "Hypothesis exploration" metric, concern the analysis of alternative explanations and possibilities. Here provenance is only marginally useful, since it can support comparative reasoning about *existing* alternatives but cannot support the automatic *generation* of alternative hypotheses. This raises a high-impact research question: how to automatically generate alternative hypotheses that plausibly explain the data? Case-based reasoning (CBR) over provenance graphs may address this research challenge, but this would not address the "cold-start" version of the problem, which likely calls for domain-specific reasoning and/or common-sense knowledge.

According to standard six: "Products should state assumptions explicitly when they serve as the linchpin of an argument or when they bridge key information gaps. Products should explain the implications for judgments if assumptions prove to be incorrect. Products also should, as appropriate, identify indicators that, if detected, would alter judgments." This is another aspect of rigor where provenance has a good start, allowing users to interactively

discover the downstream effects of assumptions proving incorrect. It is poised to take a further step in this direction – automatically searching for high-impact, linchpin assumptions (and evidentiary linchpins in general) – to present an explicit ordered list to users. The last element – identifying indicators that would alter judgments were they detected – is a natural extension of Project7 for indicators that are already represented in the provenance, but difficult for hypothetical (missing) indicators, short of perhaps incorporating CBR to use previous cases to hypothesize relevant indicators. Related to this standard, Zelik et al.'s "Sensitivity Analysis" attribute of rigor calls for evaluating the strength of an assessment given variations in source reliability and uncertainty. Project7 already supports human-led sensitivity analysis and is well-positioned to automate the search for sensitivities and quantify them. A relevant research question is: what are useful measures of sensitivity, in general or with respect to certain types of data? A sensitivity measure might have to identify sensitivities that an expert analyst or journalist would find relevant, and also express the sensitivity to the lay public, to improve data integrity literacy (more below).

Standard eight is noteworthy: "Analytic products should apply expertise and logic to make the most accurate judgments and assessments possible." Provenance does not vet the logic of an argument, nor can it verify that relevant expertise has been captured in the argument. A provenance graph instead provides a record of argumentation, which can act as scaffolding for human users to know where to investigate for logical soundness. This seems out of reach for the machine without a much richer ontology and more pervasive reference-resolution capabilities.

Standard nine is about incorporating visual information where appropriate. Project7 adheres to this standard by providing an interactive view of the provenance graph, as well as by linking items in the graph to other available views, e.g., imagery, an interactive document view for text/NLP, a map view for items with a geographical aspects, a timeline for temporally situated elements, and more.

Zelik et al.'s "Information Validation" metric is about actively and systematically vetting collected data with multiple independent, credible sources, and seeking data with convergent evidence. Project7 is able to simultaneously search multiple structured data stores (e.g., Wikidata, DBpedia, Open Street Maps), and is equipped with NLP tools that help synthesize their results, e.g. named entity recognition. It offers a canvas for constructing a concept web that integrates pieces of gathered evidence. These features can help draw out a convergence of evidence where it exists, but their coordination must be performed by a human operator. More sophisticated reasoning may enable a more proactive role for the machine-as-data-validator.

Also of note is Zelik et al.'s "Explanation Critiquing" metric of analytical rigor, which has to do with seeking feedback on an entire analysis. One positive indicator is the use of "devil's advocacy" to challenge hypotheses and explanations. A multi-user environment with support for differing appraisals like Project7 is a useful tool in this regard. Better would be automated or semi-automated devil's advocacy, which points to an interesting research question: Given a provenance graph, how can rigorous counter arguments be automatically gathered from the public sphere?

Building Public Resistance: Expectations, Literacy, and Standards

We believe an essential pillar of resiliency to information manipulation is building up widespread immunity among the public. This endeavor involves establishing higher expectations

for trustworthy information, new forms of literacy, and standards and practices that have the potential to gain traction. Below we outline ideas that connect these directions to our discussion of provenance and rigor above.

First, there is a potential link between advancing the trustworthiness of information by supporting analytical rigor and raising the public's expectations about information integrity. Empowering media organizations, intelligence analysts, and other information producers to output more reliable information may gradually build public trust in the institutions that exercise this power. In short, by supporting rigorous reporting we can raise the bar for information integrity. This may in turn make information consumers less forgiving of outlets that produce less reliable or manipulated information. Another way to drive expectations for information integrity would be to make available a public-facing interface for analyzing provenance. Imagine if authors – equipped with a provenance management interface – were able to export the full provenance for their reporting either as metadata embedded in the article or as a linked data file. The savvy information consumer might then load the article into the online provenance interface of their choice, allowing them to probe deep analytical questions and measure the information for rigor, ultimately increasing trust in information that is worth trusting.

A related pursuit is fostering public literacy about information integrity. In previous sections, we have touched on a relatively simple provenance ontology, quantitative measures of uncertainty (confidence and likelihood), and concepts from analytic rigor such as linchpin assumptions, judgment sensitivity, and hypothesis exploration. All of these would be useful to inject into the common parlance of information exchange, and the simplest way to do so is from the supply side. To start to build expectations, analysts and reporters could provide meta information to summarize the provenance, certainty, and analytical process behind the assessments that they produce. Provenance-based analytics can help by facilitating or even automating this summarization. A related subproblem is the generation of natural language descriptions of provenance graphs, which was achieved by the PROVglish architecture (Richardson & Moreau 2016). Another related research question would be how to control and automatically tailor the level of detail in the summarization. For the intelligence community this concern relates to clearances, "tear lines", and (for journalists as well) protecting sources, but it is also a relevant concern for improving public literacy about information integrity, since the sharing context might inform what level of detail is palatable for the consumer. For example, a tweet might call for a more brief summary than a feature article. Different topics may also call for different forms of sharing.

It would also be helpful to establish standards to guide expectations and ground this new literacy. A common system of interpretable tags to characterize provenance structure, certainty, assumptions, sensitivities, alternatives, etc., would make our new information-integrity literacy easier to convey, easier to use for quick comparisons and sorting, more salient, and more trendy. It would also be worth trying to establish a standard set of metrics to characterize information integrity, e.g. a measure of an assessment's sensitivity as described in the previous section. It is important that these information integrity tags and metrics are as simple as possible while capturing a useful level of detail, to give them the best chance of widespread use.

Finally, provenance technologies could improve the evolving practice of involving the public in the production of journalism, by providing structure and managing risk. Crowdsourced journalism, participatory journalism, citizen journalism and grassroots journalism have intersecting definitions, but their increasing relevance makes clear that individuals – even those outside of journalistic institutions – have more opportunities to directly participate in the

production of information that is received with higher legitimacy. These practices are not only vehicles for public literacy, but pathways by which literacy feeds back into the trustworthiness of information, making them a compelling topic of research. Crowdsourced investigations have proven powerful and unwieldy, marked by life-saving successes as well as unjust cases of misidentification and vigilantism (Venkatagiri, 2021). We can see some potential benefits of crowdsourced journalism in metrics of analytic rigor discussed above – namely hypothesis exploration, information validation, and explanation critique. Indeed, Aitamurto (2019) identified ways in which crowdsourced reporting can benefit the journalistic norms of accuracy, objectivity, and transparency. However, that paper also contained a discussion of its risks to the norms of accuracy, objectivity, and autonomy, in which a theme was the lack of structure, both in the process – which works better when led by the journalist – and in the crowdsourced information itself – which has the potential to overwhelm in unstructured form. Collaborative frameworks based on provenance ontologies could deliver the needed structure on both fronts and should be studied in this context.

Conclusions

We have targeted two interrelated questions that are critical for information integrity in today's ecosystem: (1) How can we advance the trustworthiness of information? (2) How can we boost the informational immune systems of the public? Provenance – data that describes the origins of information – is highly relevant to both questions. SIFT's formal treatment of provenance as a dependency graph has been useful for computing answers to questions of impact. We illustrated its utility by describing its role in Project7, an experimental intelligence analysis workspace that also has potential relevance to journalism and the public square. Highlighting the relevance of provenance-based analytics to tradecraft standards and rigor metrics elucidated where these tools are already innovative in advancing the trustworthiness of information and where they have potential to push it further. We finished by discussing potential benefits of these technologies for public resilience, namely higher expectations, better literacy, new standards, and collaborative frameworks for participatory journalism based on provenance.

References

Aitamurto, T. (2019). Crowdsourcing in journalism. In *Oxford Research Encyclopedia of Communication*.

Curry, A. L., & Stroud, N. J. (2021). The effects of journalistic transparency on credibility assessments and engagement intentions. *Journalism*, 22(4), 901-918.

Diakopoulos, N., & Koliska, M. (2017). Algorithmic transparency in the news media. *Digital journalism*, 5(7), 809-828.

Forbus, K. D., & De Kleer, J. (1993). *Building problem solvers* (Vol. 1). MIT press.

Friedman, S., Rye, J., LaVergne, D., Thomsen, D., Allen, M., & Tunis, K. (2020). Provenance-based interpretation of multi-agent information analysis. *Proceedings of TaPP*.

Friedman, S. E., Rye, J., McLure, M., Wauck, H. C., Patel, P., Wheelock, R., Valovage, M., Johnston, S. & Miller, C. (2021, October). Provenance as a Substrate for Human Sensemaking and Explanation of Machine Collaborators. In *2021 IEEE International Conference on Systems, Man, and Cybernetics* (SMC) (pp. 1014-1019). IEEE.

Guillory, J. J., & Geraci, L. (2013). Correcting erroneous inferences in memory: The role of source credibility. *Journal of Applied Research in Memory and Cognition*, 2(4), 201-209.

Karlsson, M. (2020) Dispersing the Opacity of Transparency in Journalism on the Appeal of Different Forms of Transparency to the General Public, *Journalism Studies*, 21:13, 1795-1814, DOI: 10.1080/1461670X.2020.1790028

Karlsson, M., & Clerwall, C. (2018). Transparency to the Rescue? Evaluating citizens' views on transparency tools in journalism. *Journalism Studies*, 19(13), 1923-1933.

Koliska, M., & Chadha, K. (2016). Digitally outsourced: The limitations of computer-mediated transparency. *Journal of Media Ethics*, 31(1), 51-62.

Lasica, J. D. (2004). Transparency begets trust in the ever-expanding blogosphere. *Online Journalism Review*, 12.

Lui, L., & Standing, L. (1989). Communicator credibility: Trustworthiness defeats expertness. *Social Behavior & Personality: an international journal*, 17(2).

McGinnies, E., & Ward, C. D. (1980). Better liked than right: Trustworthiness and expertise as factors in credibility. *Personality and Social Psychology Bulletin*, 6(3), 467-472.

Nelson, J. L., & Kim, S. J. (2021). Improve trust, increase loyalty? Analyzing the relationship between news credibility and consumption. *Journalism Practice*, 15(3), 348-365.

Office of the Director of National Intelligence. (2021). "Intelligence Community Directives," https://www.dni.gov/index.php/whatwe-do/ic-related-menus/ic-related-links/intelligence-communitydirectives, 2021.

Peifer, J. T., & Meisinger, J. (2021). The value of explaining the process: How journalistic transparency and perceptions of news media importance can (sometimes) foster message credibility and engagement intentions. *Journalism & Mass Communication* Quarterly, 98(3), 828-853.

Richardson, D. P., & Moreau, L. (2016). Towards the domain agnostic generation of natural language explanations from provenance graphs for casual users. In *International Provenance and Annotation Workshop* (pp. 95-106). Springer, Cham.

Tandoc Jr, E. C., & Thomas, R. J. (2017). Readers value objectivity over transparency. *Newspaper research journal*, 38(1), 32-45.

Venkatagiri, S., Gautam, A., & Luther, K. (2021). Crowdsolve: Managing tensions in an expert-led crowdsourced investigation. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), 1-30.

Zelik, D. J., Patterson, E. S., & Woods, D. D. (2010). Measuring attributes of rigor in information analysis. *Macrocognition metrics and scenarios: Design and evaluation for real-world teams*, 65-83.