Trust in AI in Under-resourced Environments: Lessons from Local Journalism

MARIANNE AUBIN LE QUÉRÉ and MAURICE JAKESCH, Cornell Tech, USA

What does it take for AI system to be trustworthy in under-resourced environments? As local newsrooms in the U.S. grapple with diminishing advertising revenue and decreased employment, a recent wave of philanthropic funds has promoted "AI in Local News" initiatives to champion a sustainable future for local news. Against the backdrop of economic and societal challenges facing local news organizations in the U.S. today, we identify four configurations for how local journalists use AI-related tools: *self-authored local AI tools*, *external AI tools to analyze local data*, *pre-packaged AI tools for big local data*, and *automated parachute journalism*. We then discuss the hurdles to trustworthiness and adoption to such tools considering the unique constraints of an under-resourced context. We argue that, when AI is implemented in contexts like local journalism, several aspects of trustworthy AI need to be reframed and reconsidered.

CCS Concepts: • **Human-centered computing → HCI theory, concepts and models**.

Additional Key Words and Phrases: AI-human teams, AI, trust, journalism, local news

ACM Reference Format:
Marianne Aubin Le Quéré and Maurice Jakesch. 2022. Trust in AI in Under-resourced Environments: Lessons from Local Journalism. In *CHI '22: Workshop on Trust and Reliance in AI-Human Teams.* ACM, New York, NY, USA, 9 pages.

1 INTRODUCTION

In the last twenty years, many local newspapers in the U.S. have become financially unsustainable. Compared to its 2005 peak, the estimated revenue of U.S. newspapers has fallen by over 80% [10]. This general trend is further accelerated as struggling local businesses stopped advertising during the COVID-19 pandemic [2]. Accompanying the loss of newspapers comes a decrease in the workforce of trained journalists: over 2,100 newspapers have disappeared in the last 15 years, and the total number of American newsroom employees decreased by 44% between 2010 to 2020 [1, 10]. Economic changes also affect the content that the remaining local newsrooms produce: the number of investigative pieces written by local newspapers is declining [40], once-reputable outlets are becoming "ghosts" of their former selves [1], and local newspapers that have been bought by larger conglomerates recycle content [39].

In a bid to restore the feasibility and sustainability of local newsroom, philanthropic initiatives have sought to build capacities at the intersection of local reporting and technology. Google and Facebook both provide infrastructure for local newsrooms and fund local journalism projects that leverage AI [19, 37]. Similarly, the Knight Foundation announced an "AI and Local News" initiative that provides funds to develop and advance projects that use AI to address the needs of news organizations [11].

However, in the fraught economic climate local newsrooms operate in, it is yet unclear if local newsrooms will be able to make productive use of AI systems, and if they do, whether these tools will help them become more sustainable. In this position paper, we contribute 1) a review of use

TRAIT '22, April 30, 2022, New Orleans, LA, USA
ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

cases for AI systems used in local newsrooms today, 2) implications for trustworthy AI in local newsrooms, 3) additional constraints raised by the under-resourced context that local newsrooms operate under. We argue that, when AI is implemented in under-resourced contexts, certain aspects of trustworthy AI need to be reframed and reconsidered.

2 AI FOR LOCAL STORIES

Artificial intelligence (AI) technologies use "the training of a machine to learn from data, recognize patterns, and make subsequent judgments, with little to no human intervention" [9]. In the context of news practice, the concept of AI refers to "a collection of ideas, technologies, and techniques that relate to a computer system's capacity to perform tasks normally requiring human intelligence" [24]. The specific techniques used in newsrooms vary, but often include machine learning classification or recommendation tasks, statistical data techniques, and natural language processing approaches.

While scholars and practitioners have speculated on how technology might shape our news processes (e.g. [9, 13]), the exact prevalence of AI adoption in newsrooms today seems limited and remains challenging to estimate. Many large newsrooms have integrated AI across their processes, ranging from applications like the Reuters *Tracer*, which monitors social media for breaking news events [28], to the the Associated Press's over 3,000 auto-generated earnings articles produced each quarter [41]. A recent study led by the Knight Foundation of 130 AI deployments in journalism found that 39% of the systems helped with news gathering, 15% with automatic story generation, and 13% supported news production [25]. However, the usage of AI in newsrooms is not distributed equally: the same study found that only 10% of these AI systems were used in local newsrooms.

In this section, we explore how local journalists engage with AI tools. To identify real-world examples of AI in local newsrooms, we used lists of *AI in journalism* projects compiled by media organizations and foundations, focusing on examples that were specific to local news [22, 24, 25]. We included additional examples from a literature review as appropriate. We categorize the larger list of use cases for AI in local news into four themes, which represent different levels of journalistic autonomy over the AI technology. The themes reflect the breadth of AI systems used, and showcase the tradeoffs journalists weigh between automation and efficiency, and editorial autonomy. We highlight example projects that we found most illustrative and representative for use cases in their category.

Since the field of AI in local journalism is in its infancy, many of the deployments do not match the scale and sophistication of AI systems found in other application contexts. Furthermore, some of the projects themselves are exploratory in nature. Rather than trying to create a comprehensive and in-depth analysis of what AI can do for local journalism, we provide a list of AI tools used in local journalism to guide the subsequent discussion of what it would mean for AI systems to be trustworthy in this context. We focus on tools that help the process of journalistic newsmaking only. We do not consider other AI systems that interact with news organizations, for example dynamic paywall models that have been leveraged to increase subscriptions.

2.1 Self-authored local AI tools

AI systems and tools can help drive specific local investigations and keep communities informed. Journalists can leverage AI to write stories by using locally available data, run a model or code an AI system, and publish the resulting data or output. These types of investigations require employees that are trained in machine learning, programming and/or statistical techniques. The projects below successfully tackle local issues using AI. However, they took teams of experienced data journalists and statisticians multiple months to execute and are not viable pathways for the current economic landscape for local news.

2.1.1 Los Angeles Times' QuakeBot. Employees at the Los Angeles Times identified a need to rapidly inform their communities of imminent dangers presented by frequent earthquakes. QuakeBot is an example of a human-in-the-loop AI system. The program monitors earthquake notices from the U.S. Geological Survey, and, if the magnitude exceeds 1.0, will draft an automated article. The article is then sent to humans at the LA Times for review, who may make edits to the article, choose not to publish (if the earthquake is deemed minor), or add additional detail (especially if the earthquake is of a high magnitude or impact) [38]. In January 2022 alone, QuakeBot published 10 stories to the LA Times' website, giving us an estimate for how much journalistic time has been saved. Although the application of this tool is local to California, the definition of the Los Angeles Times as a local newsroom is debatable, as they have over 1,000 employees, and are thus more likely to be able to justify the cost of such a project.

2.1.2 The Georgia Legislative Navigator. The Georgia Legislative Navigator is a website created by the Atlanta Journal-Constitution in 2015 aimed at helping "political junkies" follow along with developments at the Georgia General Assembly. The Navigator uses state-level APIs to retrieve information about members of the assembly, the status of bills, and campaign finance and lobbying data. A key feature of the website is the "Predict-a-bill" model, which uses past bill outcomes and metadata to predict and surface the likelihood that a bill will pass. The creation of this tool was a collaboration between data journalists at the Atlanta-Journal Constitution and academics at the University of Georgia [12, 30]. During its initial forty-day launch, the page received 20,000 unique views, which the creator deemed "not huge, but respectable' given the niche audience' [30].

2.2 Using external AI tools on local data

The second way that local newsrooms can leverage AI in their newsmaking process is to use their own data with AI tools that others create and maintain. Such projects leverage local journalists' own data, but abstract out the analysis or algorithmic work. The typical task for the algorithm is targeted and modular and could have likely (or would have once) been done manually, but can now be accelerated by artificial intelligence. While these types of systems still need significant investment on the part of the journalists to acquire and fit the data to a format suited to the third-party tool, these methods to do not require an in-depth understanding of machine learning or statistical methods to extract value from data. Such tools can be particularly helpful for local journalists to gain access to locally specific, but vast and unstructured data. However, tensions arise around the loss of control of the data and the algorithms that shape news.

2.2.1 Overview and the WRAL News report on food stamps. In 2013, Tyler Dukes, a reporter at WRAL News, a Raleigh, North Carolina news network, used a visual document clustering tool to uncover and report on a computer bug in a newly deployed social services system [15]. People who relied on the system, such as food stamp recipients, had complained their benefits were months overdue, while officials denied any system issues. The only information WRAL News was able to obtain came in the form of over 4,500 unordered, printed emails between users of the system and health officials. To analyze these, Tyler Dukes used a system called Overview, which was developed by the Associated Press with support from the Knight Foundation [8]. Overview helped Dukes automatically cluster emails by topic, allowing the reporter to tag relevant correspondences and ignore irrelevant clusters as appropriate. One of these clusters consisted of hundreds of emails of directors noting malfunctions in the system and being denied help at the state-level, which became the main WRAL News story. Dukes wrote a blog post to describe his experience using Overview, which describes both the benefits and downsides of using such "smart" tools for reporting. On the one hand, Dukes had to painstakingly digitize and correctly format thousands of pages to fit the systems' specifications. Yet, Dukes himself noted that using Overview was "far more effective than devoting the time and resources to do everything by hand" [16]. Since 2014, Overview is no longer in operation.

2.2.2 KPCC-LAist and Hearken create tools to parse COVID-19 questions. In March 2020, KPCC-LAist, a local newsroom devoted to the LA area, was overrun with thousands of questions from the community about COVID-19. While the reporters began by manually triaging and answering questions, they soon launched a collaboration with their tech partner Hearken and the machine learning team at Quartz to help them sort questions into thematic buckets [21]. Through sorting questions, clustering topics, and continually improving, the machine learning tool Hearken developed has been helping KPCC-LAist meet community information demands since [7]. Hearken made onboarding to an AI system easy for KPCC-LAist through the use of existing channels. The journalists were already using Hearken as a tool to collect reader questions, and Slack as a way to access the questions from Hearken. In this example, Hearken was able to add AI capabilities to help manage unexpected demand of questions, without much additional burdens on the journalists.

2.3 Pre-packaged AI tools for big local data

A third configuration of AI tools that support local journalism is the use of large-scale systems that make structured data available at scale, and provide tools to analyze and present data at a local level. The key behind these projects is a dual acknowledgement that obtaining and formatting local data can be just as challenging as its analysis. Although most of these systems do not use state-of-the-art AI techniques, they include complex statistical techniques or in some cases lightweight AI tools, and have the potential to incorporate more complex AI systems in the future. We review successful examples of such data aggregation and analysis efforts.

2.3.1 Stanford Open Policing Project and Big Local News. The Stanford Open Policing Project is an ongoing effort to collect, standardize, and analyze data about the over 50,000 traffic stops that occur in the U.S. each day [35]. The researchers that began this project requested data from police departments in all fifty states in the U.S., and received data in dozens of different file formats with inconsistent variable names. After unifying the data, they analyzed the data at scale to identify overall trends in discriminatory police stops [33]. As a next step, the researchers made this data publicly available, and put together a set of tools and trainings to help local journalists analyze data about their area, and tell effective community stories. This dataset, a close collaboration between researchers and journalists, and the creation of simple analysis tools allowed the LA Times to become one of multiple local outlets to conduct their own, lighter-weight regression analysis of police stops in a smaller area [34]. The Stanford Big Local News initiative replicates such efforts on other datasets [32].

2.3.2 MuckRock Foundation and uncounted COVID-19 fatalities. A recent attempt to combine large-scale data for local analysis is that of journalist-targeted non-profit MuckRock. MuckRock specializes in making Freedom-of-Information-Acts (FOIAs) easy to file and the resulting data easy to analyze, report on, and share with others. The MuckRock Foundation is also exploring explicitly AI approaches, such as the topical grouping of retrieved documents. As an example, the MuckRock Foundation recently paired up with journalists from USA Today and researchers at Boston University to release all the data, models, and software needed to analyze uncounted COVID-19 deaths. Like the Stanford Open Policing Project, they first published a piece with USA Today that tackles the data at a national level [27]. Next, they held a training and released how-to manuals with the necessary data to encourage local journalists to engage in their own, more pertinent analyses [6, 26]. It remains to be seen whether local journalists will take up the call to action.

Use Case	Examples	AI System Provider	Description of AI system	Goal of AI System	Local Context
Self-authored local AI tools	Quakebot, Georgia Navigator	In-house journalist or technologist	AI tool that tackles a specific problem	To automate contextual, menial tasks to free up investigative time	System is locally specific
Using external AI tools on local data	WRAL food stamps report KPCC-LAist COVID-19 questions	External entity, can be collaborative	Targeted, modular AI tool	To automate or ease part of the investigative process	Story is locally specific
Pre-packaged AI tools for big local data	Big Local News, MuckRock	External entity, can be collaborative	Large-scale AI tool that can be used by local journalists	To extract local AI insights with limited technical knowledge	Story can be locally specific
Automated parachute journalism	RADAR, Metric Media	External entity, mostly not collaborative	Fully automated local news production from larger databases	To publish local news with minimal human involvement	Story is about local events, no local context

Table 1. Overview of the four main use cases for AI systems in local newsmaking.

2.4 Automated parachute journalism

The last form of AI in local storytelling is fully automated, data-seeded local journalism. The companies that create these AI systems usually operate with data at a national or larger level, and use automation to mass-produce and distribute local stories, often without involving the local journalists or communities. Instead of supplementing existing newsroom practices, these AI systems compete with them and may replace them. We call these initiatives "automated parachute journalism," based on the idea of "parachute journalism," the practice of sending a journalist to an unfamiliar area to report on a community they have no prior context on [42]. "Parachute pieces" are often critiqued as one-dimensional or skewed portrayals of a community the journalist never gained access to, while others highlight that these pieces can serve as the only form of reporting available to an underserved community [18, 29]. Although scholars of computational journalism generally believe that quality journalism will not be replaced by AI systems [13], the instability of the local news ecosystem has created a market gap that for-profit companies are eager to fill.

2.4.1 RADAR. RADAR (Reporters and Data and Robots) is a UK-based company that calls themselves the "world's only automated local news agency" [31]. RADAR use openly available government datasets to gather data, then assign someone on their staff to manually add national context and kickstart the automation to make the stories relevant for local communities. With five data reporters, RADAR has produced over 400,000 local stories in the last three years. RADAR touts local reporters' abilities to take their stories and edit them to add further local relevance, but their main selling point for newsrooms is pure automation: one of their clients, JPI media, explained that they run 95% of stories as-is [14].

2.4.2 Metric Media LLC. Metric Media LLC, a network of over 1,200 local news websites, takes automated local journalism to a questionable extreme. Tampone, the founder of Metric Media, claims that his goal is to use proprietary software tools and public record collection systems "to rebuild and democratize community news across the country" by building thousands of local news websites [20]. Many of Metric Media's websites feature templated, automated, or low-quality local stories. Using a mix of methods including IRS records and Facebook and Google ad libraries, the Tow Center for Journalisma at Columbia uncovered that Metric Media was funded by conservative activists and non-profits, with a goal to distribute politically persuasive content across the country [5]. Metric Media is an example of AI tools being used to pollute the local media space, further eroding public trust in local media with low-quality content, and siphoning off valuable advertising revenue.

3 TRUSTWORTHY AI IN THE CONTEXT OF LOCAL JOURNALISM

In this section, we evaluate how the four use cases for AI in local newsrooms summarized in Table 1 complicate the notions of trustworthy AI, and consider additional constraints that arise specifically in under-resourced environments. Journalist's trust in AI tools and systems can be considered a special case of human-technology trust, or human-AI trust. Human-AI trust differs from interpersonal trust in that the trustee consists of two distinct entities [36]: the AI system, and the AI system provider. Some even include interpersonal trust in the AI developers in the framework of human-AI trust [23]. In our context, considering these entities as separate is crucial, as the AI systems and the AI system providers and maintainers have different goals, which in turn impact the trust placed upon the systems by local journalists. The European Ethics committee has released *Ethics Guidelines for Trustworthy AI*, with seven assessment areas for AI that is lawful, ethical, and robust [17]. We explore each of these in turn in the context of AI integration to local newsrooms.

(1) **Human agency and oversight:** AI systems in local newsrooms can either increase or decrease human agency, autonomy, and oversight. In the case of *self-authored local AI tools*, journalistic autonomy is high, as these tools are developed with the specific local context in mind, and can be used to tackle menial or repetitive tasks. For example, the LA Times specifically identified that they were producing duplicate content about earthquakes, which they preferred to automate. On the contrary, *automated parachute journalism* diminishes local journalistic agency, as the journalists do not have control over the system, even if the system has consequences on the number of readers or the trust audiences place in journalists.
(2) **Robustness and safety:** AI systems in local newsrooms must be robust and safe to limit errors in reporting from reaching audiences. Errors in reporting can have serious consequences, such as further eroding trust in local news, impacting voter behaviour, or causing panic. The impact of these errors may be heightened in local settings, where a newsroom is often the only authority reporting on local developments [1]. Minimizing error seems to be conducive to human-in-the-loop AI systems, such as the document clustering example at WRAL News. Many outside organizations (such as MuckRock, Facebook, and Google) also invest in minimizing local newsroom errors, and provide extensive documentation and training on reporting with AI tools. *Automated parachute journalism* systems have few ways to reduce errors, as they distribute too many stories for a human-in-the-loop approach, and do not work with content producers with local knowledge.
(3) **Privacy and data governance:** Ownership of local data is a recurrent theme among many of the use cases we have explored. In particular, most AI use cases in local newsrooms require sharing local data with an outside system or entity. The key distinction between the *using external AI tools on local data* and *pre-packaged AI tools for big local data* use cases is whether the data is retrieved and managed directly by the local journalist, or aggregated from some outside entity. Some AI systems, such as the *pre-packaged AI tools for big local data* also allow journalists to request public records and make them available for others to work with, thus decreasing the cost associated with FOIA requests and formatting data. This use case highlights that the sharing of local data can also be an empowering act for local journalists.
(4) **Transparency:** AI systems in local newsrooms must be transparent to local journalists and local audiences alike. For the use cases that involve the use of outside tools, this ideal can be difficult to achieve, and relies on reliable documentation and communication channels with the provider of the technology. When *using external AI tools on local data* and *pre-packaged AI tools for big local data*, the journalist must therefore take additional steps to understand the AI system, then explain the system to the audience. Algorithmic transparency may be

a smaller problem if the AI system is being used primarily for a journalism-facing process, such as news sourcing, and a larger problem if the system output is directly visible to readers, such as in automated content generation.

(5) **Diversity, non-discrimination, and fairness:** Due consideration of the local context is crucial when including AI tools in reporting. Local journalists often have a deep understanding of dynamics within their local communities, and can invoke sources from diverse backgrounds. Diversity, non-discrimination, and fairness will look different for each community, and local journalists have a responsibility to ensure that AI tools will be culturally sensitive. Local journalists can use their local context to make national discussions more pertinent for their communities, for example in the *pre-packaged AI tools for big local data* use case. Nonetheless, journalism itself is in a moment of reckoning, and diversity among local newsroom staff lags behind [4], so journalists must continue to actively engage in this area to push for change.

(6) **Societal and environmental well-being:** An active, high-quality local press keeps local powers accountable and reduces polarization and corruption [3]. In the best-case scenario, the inclusion of AI tools in local reporting allows journalists to either automate menial tasks and focus on investigative reporting, or make their investigative reporting more robust or impactful. However, in the *automated parachute journalism* use case, the use of AI tools to automate local journalism without local context may have harmful impacts on the well-being of the local community, as these types of reporting create competition for local journalists, thus harming journalistic processes that build healthy democracies.

(7) **Accountability** AI systems in local newsrooms must be accountable to the audience, who are most likely to be impacted by errors. Most newsrooms already have some systems for accountability to their audience, such as issuing corrections. However, the bearer of accountability for the system will be different based on the use case. In the *self-authored local AI tools* use case, the local newspaper themselves is continuously accountable for the system. For the *using external AI tools on local data* and *pre-packaged AI tools for big local data* use cases, the journalist and the system provider must share accountability to the audience. In the *automated parachute journalism* use case, accountability falls largely on the system provider.

While each of the seven ethical guidelines we discuss has implications for AI tools in the context of local newsrooms, in practice some are overshadowed by the constraints of working in an under-resourced environment. Table 1 illustrates how the various use cases for AI in local newsrooms can be conceptualized as a series of tradeoffs between cost of adoption, agency, goals of the system, and preservation of local context. In addition to the seven factors above, trustworthy AI systems need to be compatible with the economic situation of local journalists and complex stakeholder and power dynamics.

The economic circumstance of a newsroom is a major driving factor that influences the dimensions of trustworthy AI. Human agency and oversight is maximized in the *self-authored local AI tools*, but these custom solutions are less available to local newsrooms with fewer resources. Similarly, because local journalists do not always have the resources to build AI tools in-house, they also sacrifice elements of privacy and data governance, transparency, and accountability when using outside systems. The *automated parachute journalism* AI systems in particular seem designed to decrease job positions in local journalism while consolidating revenue and power with the system authors. Building and using AI tools has been identified as "the province of highly paid, specialized practitioners" [9], and the economic constraints of local newsrooms inhibit their ability to hire AI practitioners as part of news teams. To close this resource and skills gap, often we see local journalists collaborate with outside entities, such as non-profits, academic institutions, or companies. These collaborations may similarly involve losses in local journalist agency, oversight, privacy, data governance, and accountability over AI systems. In these ways, economic constraints mean that some of the factors of trustworthy AI become less important in under-resourced teams or must be ceded.

4 CONCLUSION

In this position paper, we consider how key elements of trustworthy AI apply to the context of local newsrooms. We present the first evaluation of how AI tools are used in the local newsmaking process. We showcase these use cases in a framework that summarizes how each use case calls for a different type of AI system, created by different entities, for different goals, with different levels of local context. We then evaluate how Floridi [17]'s seven ethical assessments for trustworthy AI apply in the context of local newsrooms. We conclude that increased constraints within under-resourced teams lead journalists to deprioritize or compromise on one or more components of trustworthy AI.

REFERENCES

[1] Penelope Muse Abernathy. 2018. The Expanding News Desert. https://uncpress.org/book/9781469653242/the-expanding-news-desert/
[2] Brad Adgate. 2021. Newspapers Have Been Struggling And Then Came The Pandemic. https://www.forbes.com/sites/bradadgate/2021/08/20/newspapers-have-been-struggling-and-then-came-the-pandemic/
[3] PEN America. 2019. *Losing the News.* Technical Report. https://pen.org/local-news/
[4] Michael Barthel. 2015. In the news industry, diversity is lowest at smaller outlets. https://www.pewresearch.org/fact-tank/2015/08/04/in-the-news-industry-diversity-is-lowest-at-smaller-outlets/
[5] Priyanjana Bengani. 2021. The Metric Media network runs more than 1,200 local news sites. Here are some of the non-profits funding them. https://www.cjr.org/tow_center_reports/metric-media-lobbyists-funding.php/
[6] Dillon Bergin and Betsy Ladyzhets. 2022. How you can use 'Uncounted' CDC data to explore the COVID-19 pandemic's hidden death toll. https://www.muckrock.com/news/archives/2022/jan/06/how-to-use-uncounted-cdc-data/
[7] Jennifer Brandel. 2021. Organizing Engagement At Scale: Hearken's New Classification Features. https://medium.com/we-are-hearken/organizing-engagement-at-scale-hearkens-new-classification-features-9b97a235bc9e
[8] Matthew Brehmer, Stephen Ingram, Jonathan Stray, and Tamara Munzner. 2014. Overview: The Design, Adoption, and Analysis of a Visual Document Mining Tool for Investigative Journalists. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (Dec. 2014), 2271–2280. https://doi.org/10.1109/TVCG.2014.2346431 Conference Name: IEEE Transactions on Visualization and Computer Graphics.
[9] Meredith Broussard, Nicholas Diakopoulos, Andrea L. Guzman, Rediet Abebe, Michel Dupagne, and Ching-Hua Chuan. 2019. Artificial Intelligence and Journalism. *Journalism & Mass Communication Quarterly* 96, 3 (Sept. 2019), 673–695. https://doi.org/10.1177/1077699019859901
[10] Pew Research Center. 2021. Trends and Facts on Newspapers | State of the News Media. https://www.pewresearch.org/journalism/fact-sheet/newspapers/
[11] Paul Cheung. 2021. AI for local news: advancing business sustainability in newsrooms. https://knightfoundation.org/articles/ai-for-local-news-advancing-business-sustainability-in-newsrooms/
[12] Atlanta-Journal Constitution. 2022. Georgia Legislative Navigator. https://legislativenavigator.ajc.com/
[13] Nicholas Diakopoulos. 2019. Automating the News: How Algorithms are Rewriting the Media. https://www.hup.harvard.edu/catalog.php?isbn=9780674976986
[14] Nicholas Diakopoulos. 2019. Boosting local news with data journalism and automation. https://www.cjr.org/tow_center/diakopoulos-automation-local.php
[15] Tyler Dukes. 2013. Records: DHHS downplayed food stamp issues. https://www.wral.com/records-dhhs-downplayed-food-stamp-glitches/13173174/
[16] Tyler Dukes. 2014. Human-Assisted Reporting Gets the Story. https://source.opennews.org/articles/human-assisted-reporting/
[17] Luciano Floridi. 2019. Establishing the rules for building trustworthy AI. *Nature Machine Intelligence* 1, 6 (June 2019), 261–262. https://doi.org/10.1038/s42256-019-0055-y
[18] Elisabeth Fondren, John Maxwell Hamilton, and Meghan Menard McCune. 2019. Parachute Journalism. In *The International Encyclopedia of Journalism Studies.* John Wiley & Sons, Ltd, 1–5. https://doi.org/10.1002/9781118841570.iejs0178 _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118841570.iejs0178.

[19] Richard Gingras. 2021. Looking at the impact of the Google News Initiative. https://blog.google/outreach-initiatives/google-news-initiative/gni-impact-report-2021/

[20] Jennifer Graham. 2020. Understanding 'pink slime journalism' and what it reveals about conservatives and liberals. https://www.deseret.com/indepth/2020/9/7/21409053/understanding-pink-slime-journalism-and-what-it-reveals-about-conservatives-and-liberals

[21] Caitlin Hernandez. 2020. KPCC is finding a new role as LA's COVID-19 help desk. Here's what it's learned along the way. https://www.niemanlab.org/2020/06/kpcc-is-finding-a-new-role-as-las-covid-19-help-desk-heres-what-its-learned-along-the-way/

[22] AI Ethics Initiative. 2018. Meet the 66 finalists in the AI and the News Open Challenge. https://aiethicsinitiative.org/news/2018/12/3/meet-the-66-finalists-in-the-ai-and-the-news-open-challenge

[23] Mark Jacob. 2020. COVID-19 Accelerates Local News Trends, for Bad and Good. https://localnewsinitiative.northwestern.edu/posts/2020/04/22/local-news-pandemic/

[24] JournalismAI. 2021. *AI Journalism Starter Pack*. Technical Report. https://docs.google.com/document/d/1pWwbqPERg0bUbMHMbYYDWmFQmWJYvK8N2Dmbenp4Qu0/

[25] John Keefe, Youyou Zhou, and Jeremy Merrill. 2021. The present and potential of AI in journalism. https://knightfoundation.org/articles/the-present-and-potential-of-ai-in-journalism/

[26] Derek Kravitz. 2021. Join MuckRock and USA Today Dec. 15 for a training on COVID death undercounting. https://www.muckrock.com/news/archives/2021/dec/14/join-muckrock-and-usa-today-for-a-training-on-covi/

[27] Derek Kravitz, Ashley White, Andrew Capps, Jake Kincaid, Nada Hassanein, Dillon Bergin, and Sarah Haselhorst. 2021. Uncounted: The hidden death toll of the COVID-19 pandemic. https://www.muckrock.com/news/archives/2021/dec/23/uncounted-the-hidden-death-toll-of-the-pandemic/

[28] Xiaomo Liu, Quanzhi Li, Armineh Nourbakhsh, Rui Fang, Merine Thomas, Kajsa Anderson, Russ Kociuba, Mark Vedder, Steven Pomerville, Ramdev Wudali, Robert Martin, John Duprey, Arun Vachher, William Keenan, and Sameena Shah. 2016. Reuters Tracer: A Large Scale System of Detecting & Verifying Real-Time News Events from Twitter. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management (CIKM '16)*. Association for Computing Machinery, New York, NY, USA, 207–216. https://doi.org/10.1145/2983323.2983363

[29] Aliyu Odamah Musa and Muhammad Jameel Yusha'u. 2013. Conflict reporting and parachute journalism in Africa: A study of CNN and Al Jazeera's coverage of the Boko Haram insurgency. *Journal of Arab & Muslim Media Research* 6, 2-3 (Sept. 2013), 251–267. https://doi.org/10.1386/jammr.6.2-3.251_1

[30] Susannah Nesmith. 2015. How Legislative Navigator tracks Georgia lawmakers. https://www.cjr.org/united_states_project/atlanta_journal_constitution_legislative_navigator.php

[31] RADAR News. 2022. Radar news - AI Generated Local News Stories. https://pa.media/radar-news/

[32] Stanford Big Local News. 2022. Big Local News. https://biglocalnews.org/

[33] Emma Pierson, Camelia Simoiu, Jan Overgoor, Sam Corbett-Davies, Daniel Jenson, Amy Shoemaker, Vignesh Ramachandran, Phoebe Barghouty, Cheryl Phillips, Ravi Shroff, and Sharad Goel. 2020. A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour* 4, 7 (July 2020), 736–745. https://doi.org/10.1038/s41562-020-0858-1 Number: 7 Publisher: Nature Publishing Group.

[34] Ben Poston and Cindy Chang. 2019. Analysis: LAPD searches blacks and Latinos more often - Los Angeles Times. https://www.latimes.com/local/lanow/la-me-lapd-searches-20190605-story.html

[35] Stanford Open Policing Project. 2022. The Stanford Open Policing Project. https://openpolicing.stanford.edu/

[36] Keng Siau and Weiyu Wang. 2018. Building Trust in Artificial Intelligence, Machine Learning, and Robotics. *Cutter Business Technology Journal* 31 (March 2018), 47–53.

[37] Fidji Simo. 2017. Introducing the Facebook Journalism Project. https://www.facebook.com/journalismproject/introducing-facebook-journalism-project

[38] Los Angeles Times. 2019. What is the Quakebot and how does it work? https://www.latimes.com/la-me-quakebot-faq-20190517-story.html

[39] Benjamin Toff and Rasmus Kleis Nielsen. 2018. "I Just Google It": Folk Theories of Distributed Discovery. *Journal of Communication* 68, 3 (June 2018), 636–657. https://doi.org/10.1093/joc/jqy009

[40] Eray Turkel, Anish Saha, Rhett Carson Owen, Gregory J. Martin, and Shoshana Vasserman. 2021. A method for measuring investigative journalism in local newspapers. *Proceedings of the National Academy of Sciences* 118, 30 (July 2021), e2105155118. https://doi.org/10.1073/pnas.2105155118

[41] Erin Madigan White. 2015. AP Definitive Source | Automated earnings stories multiply. https://blog.ap.org/announcements/automated-earnings-stories-multiply

[42] Sharyn Wizda. 1997. Parachute journalism. *American Journalism Review* 19, 6 (July 1997), 40–45.