

UCLA

UCLA Electronic Theses and Dissertations

Title

Algorithms and Articles: Examining the Impact of Generative Artificial Intelligence on the Production and Consumption of News

Permalink

<https://escholarship.org/uc/item/6qh0j850>

Author

Wasdahl, Alexander

Publication Date

2024

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Algorithms and Articles: Examining the Impact of Generative Artificial Intelligence on the
Production and Consumption of News

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of
Philosophy in Information Studies

by

Alexander John Wasdahl

2024

© Copyright by

Alexander John Wasdahl

2024

ABSTRACT OF THE DISSERTATION

Algorithms and Articles: Examining the Impact of Generative Artificial Intelligence on the
Production and Consumption of News

by

Alexander John Wasdahl

Doctor of Philosophy in Information Studies

University of California, Los Angeles, 2024

Professor Ramesh Srinivasan, Chair

This study assesses the impact of generative artificial intelligence (AI) on journalism in terms of news production and consumption. It advances existing research in journalism theory by examining how the rapidly advancing and contemporary field of generative artificial intelligence is being integrated into newsrooms. The study begins by reviewing relevant literature surrounding artificial intelligence and journalism, before introducing the theoretical considerations and approaches deemed best suited to approach the three research questions of the study:

RQ1: How do individuals involved in automated news production describe their experiences?

RQ2: What are the evaluative criteria used by readers in their perception of automated news content?

RQ3: How do news readers in the United States perceive automated news content and human-written news content relative to the descriptors derived from RQ2?

Through a mixed methods approach of a) semi-structured interviews with journalists and product developers describing their experiences producing content with AI and b) survey experiments evaluating and comparing perceptions of news articles generated using AI (ChatGPT), this study underscores the idea that AI-generated media is a fundamentally distinct entity from human-generated media. Thus, this entity must be reckoned with using a novel set of theoretical assumptions if we are to properly ascertain its social and cultural impacts. The study concludes by considering the future of creativity in an AI-generated world and the salience of alienation as a potential social consequence of the increased role of generative AI across industries. Ultimately, this study provides a foundation for offering a more comprehensive understanding of the multifaceted nature of news production and consumption.

The dissertation of Alexander John Wasdahl is approved.

P.J. Lamberson

Leah Lievrouw

Davide Panagia

Ramesh Srinivasan, Committee Chair

University of California, Los Angeles

2024

Table of Contents

1.	Introduction	1
1.1.	AI-Generated Text in Popular Culture	2
1.2.	Thesis Outline	4
2.	Literature Review	11
2.1.	Information Studies and Journalism	11
2.2.	Automation and AI in News Production	16
2.3.	Why is Journalism being Automated?	24
2.4.	ChatGPT and Natural Language Generation	28
2.5.	Algorithms: Representation versus Correlation	33
2.6.	Previous Studies on Automated News: Reach and Practice	38
2.7.	Research Questions	47
3.	Theoretical Considerations	48
3.1.	Overview	48
3.2.	The Algorithmic Ontology	50
3.3.	Journalism and Language	56
3.4.	Readerly Expectations and Values	60
3.5.	Epistemological and Phenomenological Assessments	65
3.6.	A ChatGPT Case Study	69
4.	Methods	78
4.1.	Purpose of The Study	78
4.2.	Research Questions	81
4.3.	Study Design	84
4.3.1.	Interview Data Collection Overview	84
4.3.2.	Survey Data Collection Overview	85
4.3.3.	Interview Participant Recruitment and Data Collection Protocol	87
4.3.4.	Survey Participant Recruitment and Data Collection Protocol	89
4.4.	Interview Analysis Overview	94

4.5.	Survey Analysis Overview	96
5.	Interview Analysis	98
5.1.	AI as a Tool	99
5.2.	Innovation	109
5.3.	Business Models	118
5.4.	Audience	125
5.5.	Trust and Transparency	132
5.6.	Guardrails	138
5.7.	The Human Element	148
5.8.	Value Alignment	157
5.9.	Summary	163
6.	Survey Analysis	166
6.1.	Phase 1	166
6.2.	Phase 2	168
6.2.1.	Exploratory Factor Analysis	170
6.2.2.	Factor Labels - Varimax Rotation	177
6.2.3.	Factor Labels - Oblimin Rotation	182
6.3.	Phase 3	188
7.	Discussion and Conclusion	197
7.1.	Considerations	198
7.1.1.	Technology, Analysis, Mythology	198
7.1.2.	Trust and Transparency	202
7.2.	Limitations	205
7.3.	Implications	208
7.3.1.	Creativity	208
7.3.2.	Alienation	210
7.3.3.	Looking Ahead	211

List of Figures and Tables

Figure 1.1: An example of a popular BONY Facebook post	3
Figure 2.1: An AI-generated author profile at Sports Illustrated (has since been deleted)	22
Figure 2.2: An example of erroneous AI output in a Gannett newspaper article	23
Figure 3.1: The ChatGPT (GPT-4) interface	71
Figure 3.2: A ChatGPT Hallucination Example	73
Figure 3.3: Another ChatGPT Hallucination Example	74
Figure 3.4: A News Article Prompt	76
Figure 3.5: LinkedIn Header for Dr. Strudley	77
Figure 4.1: Prompt for Phase 1 Survey	87
Figure 6.1: Scree Plot (Varimax Rotation)	171
Figure 6.2: Scree Plot (Oblimin Rotation)	172
Figure 6.3 Model estimates (finance topic only)	195
Figure 6.4 Model estimates (politics topic only)	196
Figure 6.5 Model estimates (tech topic only)	196
Table 2.1: Automated News Production and the Journalistic Value Chain	20
Table 4.1: Survey Phase Design Summary	90
Table 4.2: Phase 3 Article Length	94
Table 6.1: Eigenvalues and Variance Explained (Varimax Rotation)	173
Table 6.2: Eigenvalues and Variance Explained (Oblimin Rotation)	174
Table 6.3: Factor Loadings of News Stories Ratings - Varimax Rotation	175
Table 6.4: Factor Loadings of News Stories Ratings - Oblimin Rotation	180

Table 6.5: Factor Correlations - Oblimin Rotation	186
Table 6.6: Engagement Ratings (Mean, SD) by Authorship and Topic	190
Table 6.7: Substance Ratings (Mean, SD) by Authorship and Topic	190
Table 6.8: Clarity Ratings (Mean, SD) by Authorship and Topic	191
Table 6.9 Prolixity Ratings (Mean, SD) by Authorship and Topic	192
Table 6.10 Alienation Ratings (Mean, SD) by Authorship and Topic	192
Table 6.11 Complexity Ratings (Mean, SD) by Authorship and Topic	192
Table 6.12 Effort Ratings (Mean, SD) by Authorship and Topic	192
Table 6.13 Trust Ratings (Mean, SD) by Authorship and Topic	193
Table 6.14 Neutrality Ratings (Mean, SD) by Authorship and Topic	193

Acknowledgements

There are a great many people that I would like to thank here, including first and foremost my committee members: P.J. Lamberson, Leah Lievrouw, Davide Panagia, and Ramesh Srinivasan, whose continued advice and feedback throughout my doctoral program and dissertation journey has helped me immensely. I am also very thankful to my parents Dan and Colette and siblings Coral and Ivy, whose support got me through the toughest phases of this process. I would also like to thank my interviewees and survey respondents, who were able to take the time to participate in my research. Finally, I would like to thank Carly Rae Jepsen, whose 2015 magnum opus *Emotion* provided me with a comforting soundtrack of some of the greatest pop music of all time during many long nights of writing.

Vita

Alex Wasdahl's research examines the production and consumption of AI-generated news content. He has published writing that forwards different theoretical and practical approaches to generative AI ethics and policy. During his time at UCLA, he served as Head Teaching Fellow for the Data, Justice and Society Cluster in 2023-2024, serving as lead instructor for a spring 2024 course on AI and journalism. He has also held teaching assistantships in the Masters of Social Science (MaSS) and Digital Humanities programs from 2019-2023, which informed his research interests. Alex received a Dissertation Year Fellow award for the 2023-2024 academic year and also received the David Vickter Memorial and Andrew Horn Memorial scholarships as a doctoral student. Alex also assumed a variety of leadership roles on campus at UCLA, including Director of Communications for the Graduate Student Association, New Ventures Fellow at the Technology Development Group, and President of the Advanced Degree Consulting Club. He holds a B.A. in Government from Dartmouth College.

Chapter 1: Introduction

The burgeoning influence of artificial intelligence (AI) on news production, manifested in part in the form of article writing, has gained significant momentum over the past decade. Less than 10 years ago, automation in news was restricted to template-based approaches (i.e. an algorithm fills out a human-written template with structured data from a separate dataset); now people can prompt user-facing generative AI models to write freeform text (Wang and Huang, 2024; Leppanen et al., 2017; Marvin et al., 2023). This acceleration is largely powered by the advent of advanced large language models (LLMs). The release of LLMs such as OpenAI's ChatGPT and Google's Bard have garnered significant hype and media attention since late 2022; such advancements underscore the escalating relevance of generative AI across numerous publicly salient domains. This dissertation advances existing research in journalism theory by examining how the rapidly advancing and contemporary field of generative artificial intelligence is being integrated into newsrooms. Entrepreneur Magazine recently declared 2023 as the “Year of Generative AI,” and MIT Technology Review contends that generative AI is “one of the biggest upheavals in a generation.” (Heaven, 2022; Hashmi, 2023) Multiple news outlets have experimented with generative AI technology to write articles, sometimes with minimal editorial oversight (Bauder, 2023; Knibbs, 2024). As evidenced by the 2023 UCLA town hall on ChatGPT, it is crucial that institutions take the initiative to engage with the potential of this revolutionary technology (Harmon, 2023).

The recent impact of artificial intelligence, and generative AI in particular, on journalism has been transformative, redefining the scope and methodology of news production. The integration of AI tools in journalism is not just a futuristic concept but a present reality, with applications ranging from routine reporting to complex data analysis. One of the most

straightforward applications is in sports journalism, such as automating reports on high school football scores. Similarly, AI can be used to analyze vast troves of leaked documents by parsing extensive data to uncover insights more efficiently than humans. These tasks, traditionally time-consuming and labor-intensive, are now more accurate with the assistance of artificial intelligence (Beckett, 2023). Parallel to this technological evolution, the academic sphere has seen a surge in research exploring both journalist reactions to AI integration and reader perceptions of AI-generated news (Graefe and Bohlken, 2020; Henestrosa et al., 2023; de Haan et al., 2022).

The COVID-19 pandemic serves as a contemporary case study of artificial intelligence capabilities when applied to journalism. As noted by Danzon-Chambaud (2023), the outbreak saw governments and health authorities release substantial amounts of open-source data, accessible through structured datasets or APIs. This data included vital statistics like death tolls, ICU patient numbers, and incidence rates. The structured nature of this data made it conducive to automated news coverage, enabling media organizations to efficiently report on the evolving situation with a high degree of accuracy and timeliness. ChatGPT, with its advanced language generation capabilities, presents another intriguing case study. As Diakopoulos (2023) notes, while automated writing based on structured data is a long-standing practice in journalism, models like OpenAI's Generative Pre-Trained Transformer (GPT) bring a nuanced complexity to this process.

1.1 AI-Generated Text in Popular Culture

On May 2nd, 2019, the Facebook group Bots of New York was created. Bots of New York--or BONY as it is often called--is a parody of the popular blog Humans of New York, a

collection of street portraits and interviews in New York City launched by author and photographer Brandon Stanton in 2010. The parody group Bots of New York features AI-generated photographs of New Yorkers paired with AI-generated stories generated using the GPT-2 deep learning language model. The AI-generated photographs typically feature the face of a particular subject embedded on the body/profile of another subject. By operating on Humans of New York blog posts as training data, the algorithm that powers the Bots of New York page generates absurd and sometimes oddly realistic posts with hilarious results (see Fig. 1.1).

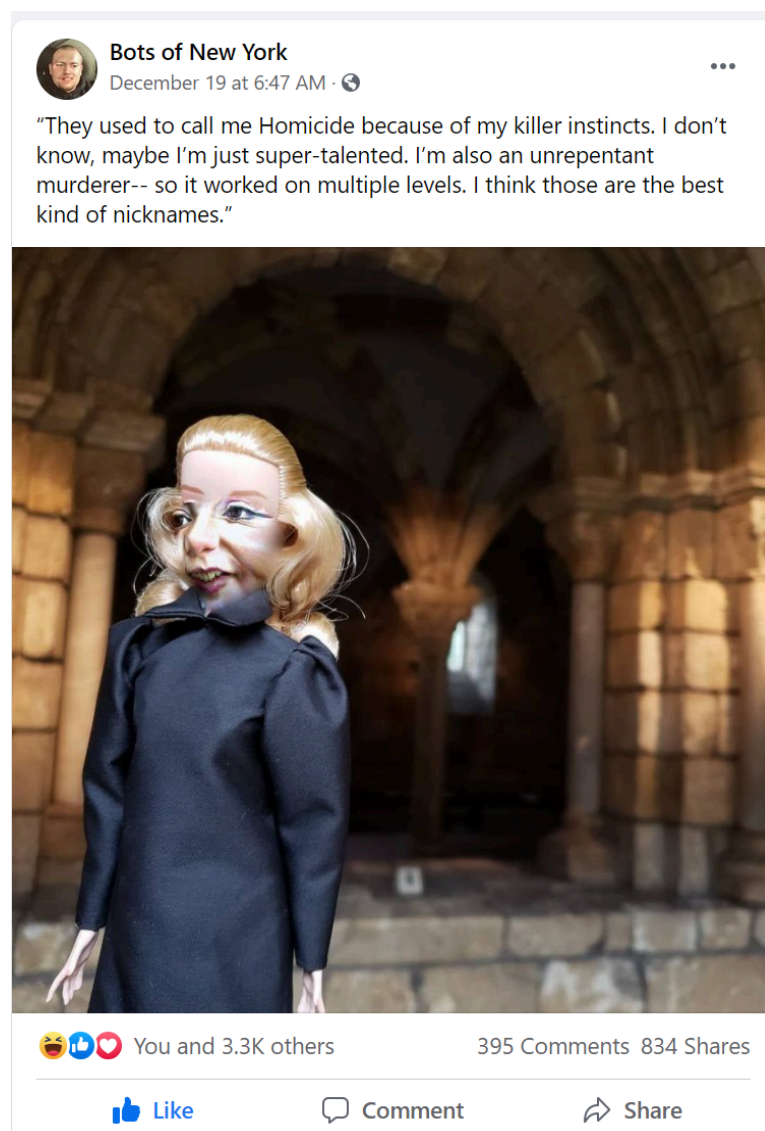


Figure 1.1: An example of a popular BONY Facebook post

The posts on the Bots of New York page that perform best with regard to likes and shares typically possess one or more of the following characteristics: 1) puns or plays on words that insinuate a command of irony on behalf of the (AI) author, 2) repetitive or otherwise nonsensical sequences that convey a sense of absurdity, or 3) portraits that result from bizarre or unusual combinations of subjects. Figure 1.1 possesses the first and third of these characteristics: it features both a witty caption and an unusual face mashup. The first sentence, “They used to call me Homicide because of my killer instincts,” contains a play on the word “killer” having two different meanings: a “killer” as in one who commits homicides and “killer instinct” as an idiom referring to an inclination to succeed at all costs. Based on this sentence, one might assume that the GPT-2 algorithm was clever enough to understand the humorous nature of this double meaning. However, this assumption belies the reality of artificial neural networks such as GPT-2.

Since 2019, the training and text-generation capabilities of these models have progressed rapidly, demonstrated by the release of GPT-3 in 2020 and GPT-4 in 2023. Most notably, OpenAI released a user-friendly chatbot, ChatGPT, built on top of their foundational LLMs in late 2022. The ChatGPT chatbot quickly went viral on social media as users shared examples of what it could do. Stories and samples included everything from planning travel agendas, crafting computer code, and even writing news articles. Within five days, the chatbot had attracted over one million users (Marr, 2023).

1.2 Thesis Outline

This dissertation explores engagement with generative artificial intelligence in the news production and consumption processes. Production in this case refers not just to writing but the

entire process of creating news from ideation to post-distribution. Consumption refers to the process of reading and engagement with news content created by generative artificial intelligence. The burgeoning influence of generative AI in journalism inevitably brings to the forefront the critical importance of understanding how both writers and readers engage with this evolving entity.

With regard to news production, the problems that algorithms pose as political entities with powers of governance are ones with which to confront the world of journalism. Elements of this confrontation include the priorities of journalistic practitioners: what they want, what they think about this new medium, and accounts of their experiences with the artificial intelligence tools that facilitate automated text generation. In addition, it is crucial to consider the extent to which the adoption of new algorithmic newsroom technologies would potentially alter the way in which they write articles, bylines, and other content. The autonomy of practitioners within the realm of journalism also warrants examination in terms of how newsrooms balance journalistic and algorithmic labor in the areas of reporting, editing, and curating. Increased awareness of the promises and pitfalls of large language models (LLMs) has led to concerns about their potentially deleterious effects in the industries in which they are being experimented. Journalism, as the canary in the coal mine for generative AI, is a prime example (Leiser, 2022). As news organizations both large and small reckon with this rapidly advancing technology, it is of paramount importance to assess how generative AI affects the experience of news workers across the journalistic value chain—not just news writing, but gathering and distribution as well.

With regard to news consumption, building upon foundational research in information studies and media theory, this study of AI-generated news is particularly timely. An evaluation of perceptions of AI-generated news content aims foremost to understand essential components of a

good AI-written news story. In order to understand this, it is crucial to assess the essential components of a good news story, and whether or not these components are the same as those of a good AI news story; and furthermore, should they be the same? Moreover, are these components attended to by receivers of this content (i.e. news readers)? This dissertation considers these factors as its foundation in its exploration of the evaluative criteria readers use to perceive automated news content. In the broader context of news research, much emphasis has been placed on distinguishing news content from other communication forms. This is based on the understanding that each content category has its unique evaluative criteria considered pivotal. The advent of AI-generated news as a novel content concept demands renewed attention toward modes of understanding reader perceptions. To make sense of how readers perceive AI-generated news, it is of paramount importance to unravel the criteria that underpin these perceptions. As the media world assesses the impact and potential of artificial intelligence in journalism, this research becomes crucial to understanding and navigating this evolving landscape.

As Beckett (2023) astutely observes, the efficiency and rapidity with which generative AI can disseminate good information can, unfortunately, be equally applied to the spread of misinformation and disinformation. This dual-edged nature of AI-generated content underscores the necessity of strategic thinking, not just in terms of content generation but also in fostering media literacy and resilience among readers. As AI tools become more prevalent in newsrooms, the focus should shift towards empowering readers to discern and critically evaluate the content they consume. This involves a comprehensive approach that goes beyond merely tracking the volume of misinformation and focuses on cultivating a more discerning and informed audience. Such efforts are essential to ensure that the advancements in AI-driven journalism serve to enhance public discourse and knowledge, rather than detract from it due to the unchecked spread

of false or misleading information. This shift towards media literacy resilience is not just a defensive strategy against bad actors but a proactive measure to uphold the integrity and trustworthiness of AI-enhanced journalism in the digital age.

News organizations' decisions to deploy generative AI tools for content creation are significantly influenced by considerations of how such content is perceived by their audience. As outlined in the Modality-Agency-Interactivity-Navigability model by Sundar et. al. (2008), there exists a complex and often conflicting set of perceptions among readers regarding automated news. On one hand, readers may exhibit a preference for human-written articles, attributing a higher level of expertise to journalists based on the authority heuristic. This preference could also stem from a desire for a sense of human connection or social presence, favoring communication with a human rather than a machine. Conversely, the machine heuristic suggests another perspective, where readers may view automated news as more objective, being free from ideological biases often associated with human journalists. These contrasting views present a unique challenge for news organizations as they weigh the economic potential of AI against reader trust (or lack thereof). To navigate these complexities, researchers in various countries have engaged in experimental studies aimed at understanding and comparing reader perceptions of automated news with that of human-written content. Such research is crucial for news organizations as they assess the feasibility and impact of integrating AI tools into their content creation processes, balancing technological advancement with audience expectations and preferences.

This dissertation ultimately seeks a comprehensive understanding of both individual experiences with automated news production and reader perceptions of automated journalism content. To accomplish this, I begin in Chapter 2 by reviewing relevant literature surrounding

artificial intelligence and journalism. Starting with the historical origins of the concept of artificial intelligence, the chapter subsequently covers the increase of datafication and the recent advent of user-facing generative AI tools in the early 2020s as the contemporary context in which AI-generated content is an object of examination and critique. I present generative AI as a specific type of artificial intelligence: traditional AI systems are primarily used to analyze data and make predictions, while generative AI goes a step further by creating new data similar to its training data. I then situate the phenomenon of AI-generated content within the theoretical boundaries of the field of Information Studies, assessing how documents relate to information and how AI-generated content serves as a new stage in the evolution of writing. Subsequently, I overview journalism as a specific case of document writing, and review existing literature that contends with AI as a dynamic disruptor of both news production and consumption. After identifying gaps in existing research—specifically a need for a reevaluation of existing metrics for evaluating news perceptions, I conclude with the research questions of the study.

In chapter 3, I introduce the theoretical considerations and approaches deemed best suited to approach said research questions of the study, which cover ontological and phenomenological dimensions of automated news content. To foreground each of the three research questions, given that they each address the phenomenon of automated news content, an ontological assessment of automated news content is key. This chapter elucidates the importance of understanding the mode of existence of an algorithm in order to understand its effects. In chapter 4, I delve into the specific methodological choices that underpin the study, first articulating the purpose of the study and the specific gaps that it fills in the literature on AI-generated news content. Next, I present the research questions and study design, including considerations regarding participants, platforms, and instruments. In this chapter I also provide justifications for

the methodological choices, including the importance of a mixed methods study for this particular topic and the complementary value of both interview and survey data. Following this, an account of procedures in both the interview and survey stages of the study is provided, including data collection and analysis protocols.

In chapter 5, I analyze the semi-structured interviews that were conducted to address the corresponding research question. The chapter is organized into coded themes that were identified during the analysis process. For each theme, responses by interview participants that relate to or connect to the theme are included and discussed. Within each of the themes, sub-themes are discussed and supported with quotes. In this chapter, through accounting the experiences of interview participants, I illustrate the evolution of journalism and technology as separate fields as well as in a fundamentally intertwined set of institutions and values. The analysis also underscores, among others, a common thread among participants experiencing AI-generated content differently than human-generated content, both in newsrooms and as objects of research. In chapter 6, I analyze the results of a three-phase survey experiment that was conducted to answer the corresponding research questions, and the chapter is organized according to each phase of the experiment. Results of each phase are analyzed and situated within the context of the relevant research questions.

In chapter 7, I reflect on the breadth of analyses covered in the study: the theoretical frameworks and ChatGPT case study, the semi-structured interview responses and theme extraction, and each phase of the survey experiment. In this final analysis, this chapter addresses some of the considerations raised by each phase of the study, with an eye toward a critique of the work itself—what it has done and what it has not done. Subsequently, this chapter addresses some

limitations of the study overall before laying out further theoretical implications, both in relation to and in extension to what the study has covered.

Chapter 2: view

2.1 Information Studies and Journalism

The increasing adoption of automated processes in news production has been catalyzed in part by the era of “Big Data,” which has brought with it a deluge of information produced by and about people. Big Data as a cultural phenomenon is based on the interplay of three forces: technology, analysis, and mythology (Boyd and Crawford, 2012). The first force, technology, involves the maximization of computer power through the accuracy of algorithms that operate on large data sets. The data sets upon which language generation algorithms operate are becoming larger with each new advancement. GPT-4 (which stands for Generative Pre-Trained Transformer-4), the most recent large language model (LLM) released to the public by OpenAI, operates on a number estimated to be in the trillions. Its previous model, GPT-3, operates on 175 billion parameters, whereas the model before that, GPT-2, only operates on 1.5 billion parameters, an increase of over 100x.

The second force, analysis, involves the identifications of patterns in data sets in order to make claims. Natural language generation models are trained on data from a massive corpus of human writing. Within this data, analysis of patterns of word sequences allow these models to accrue predictive power. The third force, mythology, refers to the conceptualization of data sets as truthful, objective providers of a higher form of intelligence and knowledge due to their ability to provide insights unattainable through other means. The application of this concept can be seen in the marketing and branding of artificial intelligence initiatives: while such technology is seen as beneficial for its ability to allow rapid access to more relevant information, it raises complex ethical questions about the public and the extent of fragmentation of the public sphere.

Ultimately, as Boyd and Crawford explain, Big Data “reframes key questions about the constitution of knowledge, the processes of research, how we should engage with information, and the nature and the categorization of reality” (ibid.).

Information--in particular, the question of “what is information”--is a question of immense scholarly concern in the field of Information Studies (IS). In seeking to answer this question, those who have contributed to this body of literature have done so in the context of any (or any combination) of the fields of science, social science, or philosophy (Furner, 2004). For example, the work of Norbert Wiener and Claude Shannon in the field of cybernetics in the 1940s offered up a narrowly scientific definition of information as a quantifiable measure of the probability of a set of bits being transmitted within a particular context (Shannon, 1948). This definition, however, does not place emphasis on the many everyday ways in which human beings interact with information through seeking it out and utilizing it. For instance, in the field of information behavior research, the term “information” is generally assumed to cover any instance where individuals interact with their environment in any such way that leaves some impression on them (Bates, 2010).

At a broad level, information represents the vast epistemic realm between fact and knowledge (Gullory, 2014). The difference between information and fact is based on value in transmission: a fact becomes information when value is added to it. A fact functions as information in certain contexts when the fact is what one wants to know in that context. Information, therefore, has a shelf life of the time that it is sought to be known, beyond which it must be stored to await its next opportunity. It is this shelf life that drives the development of our information technologies in order to optimize the speed and effectiveness of transmission (ibid.). John Seely Brown distinguishes between information and knowledge in three ways. First,

knowledge usually entails a human knower, whereas information is treated as an independent and self-sufficient entity that is associated with someone (the knower). Second, knowledge appears harder to conceptualize as a self-contained substance than information, less amenable to the ideas of shipping, receiving, and quantification than information. Third, knowledge requires more assimilation than information, entailing a certain degree of commitment on behalf of the knower (Brown, 2017). It is within this context that John Guillory proposes a definition of information as “any given datum of our cognitive experience that can be materially encoded for the purpose of transmission or storage” (2004, p. 110). In the context of writing, a definition of information that emphasizes its transmission and storage is of paramount importance.

Michael Buckland divides the meaning of information into three categories: information as knowledge--such as the stuff one gets, information as process--such as how one gets it, and information as thing--referring to the source. In this third category of “information as thing”, in which information denotes physical media, the word information is commonly used to refer to a document, a term which Buckland uses to refer to a class of objects including “bits, bytes, books, and other signifying objects” (Buckland, 2017). Within this category in which information refers to documents, the term document itself has three views that can be identified according to Buckland: a conventional/material view of documents, an instrumental view of documents, and a semiotic view of documents. The conventional view of documents conceives of documents similarly to how they are perceived in everyday life: written text records that are material, local, and generally transportable (ibid.). A news article, for example, is an instance of a document under this view because it is text written on paper (or an online text processor) and is an object with material and local properties.

Apart from the conventional view of documents, there is also both an instrumental element to a document and a semiotic element that are important to consider, particularly in the context of automation becoming a pervasive factor in many contemporary documents. The instrumental view considers anything that signifies something or acts as evidence of something to be a document. A classic example of this is Suzanne Briet's antelope analogy. A prominent 20th century documentalist, Briet sought to understand the epistemological dimension of documentation itself as well as its role in social and economic development. She defined a "document" as "any concrete or symbolic indication, preserved or recorded, for reconstructing or for proving a phenomenon, whether physical or mental" (Maack, 2014). Briet also identified several key roles of documents, chief among them being its "indexicality"--an organized relationship with other evidence grants document status to an object (Buckland, 2017). She gives the example of an antelope: running wild in the plains, an antelope is not considered to be a document. However, when it is captured and shown in a zoo it becomes a document: its status as an object of study relative to an observer (e.g., a zoo attendee) renders it so (Maack, 2004).

In the context of automated written content, it is important to consider the nature of information and documentation to understand the terminological nature of automated writing, the ways in which it represents a new stage in the evolution of writing, and how it manifests itself in the field of journalism. The act of writing is considered to be the origin of information technology. Michael Hobart and Zachary Schiffman (2000) contend that orality, while utilized as a means of passing on knowledge, values, and beliefs in ancient culture, exists not as information but as a practice whose preservation is a byproduct of repeated usage. Much like the term "information," the term "writing" can be hard to define due to the complexity of its functions. Particularly, it is important to distinguish between genres of writing in order to properly situate

automated written content within a historical context. John Guillory identifies three primary genres of writing and places them on the following epistemic axis:

literary/journalistic----informational----scholarly/scientific

At one pole, the literary/journalistic genre covers epistemic concepts of fiction and opinion, and at the other, the scholarly/scientific genre covers disciplinary forms of knowledge. Lying in between the two ends is the informational genre, within which Guillory situates the document (e.g. the form, memo, or report). His decision to situate journalism within the same genre as literary writing seems to be an interesting one, given that the purpose of news is to inform the public. However, he justifies this placement by asserting that the informational genre consists of modes of writing in which information constitutes the chief generic determinant. While information is distributed along the entire axis (such as with news in journalism and data in science), it is only the chief determinant of content within the informational genre, such as a form, memo, or report. The document is the carrier of information and therefore is the object of knowledge rather than knowledge itself (Guillory, 2004). Applying this to Briet's antelope, the knowledge represents the antelope in the wild and the form/memo/report is the antelope in the zoo. Furthermore, as Seely Brown notes, documents do not merely carry information, they help make it, structure it, and validate it (Brown, 2017).

Guillory's definitional prescription distinguishes journalism from purely informational genres, by virtue of the fact that in journalism strands of opinion and entertainment are bundled with transmitted information. He describes the epistemic positioning of the informational genre as both ephemeral and permanent, with an audience of one or none, that may be read once or

never, and is preserved or filed away. In contrast, literature aspires to eternal reading or “canonicity” and science releases knowledge from its “prison” of the written form (Guillory, 2004, p. 111). An algorithm such as one involved in news production uses input data to create output data. These two types of data represent the information, or documents, to which attention ought to be directed. Automated journalism with no human intervention acts on structured data often found in relational data sets. In the case of deep learning language models creating written text based on human prompts and writing articles with content that is cherry picked and edited by human intermediaries, the data these models act upon are the written content parameters it has been trained on or fed. The distinctions between journalism and the informational genre are less clear in the era of Big Data.

The integration of automation into the journalistic profession warrants a revisitation of Guillory’s definitional prescription. Big Data provides “destabilizing amounts of knowledge and information that lack the regulating force of philosophy” (Berry, 2011). The production of documents at enormous scale has also transformed the relationship between information, individuals, and society. How we as individuals and as a society use and understand information are part of our culture. Changes in the nature of written text thus have the potential for massive social and cultural upheaval, reinforcing the importance and timeliness of studying automated fake news as a sociocultural phenomenon.

2.2 Automation and AI in News Production

The offloading of intellectual and cognitive labor to computers is a defining feature of the “second machine age” (Diakopoulos, 2019). Whereas, in the first machine age, physical work was the primary form of labor to be automated (such as the steam engine), technological

advances (progress in computer processing, artificial intelligence, networked communication and the digitization of just about everything) have allowed for the offloading of cognitive labor--the act of judging and making decisions--to the realm of computers and algorithms. An algorithm is a series of steps that is undertaken in order to solve a particular problem or to accomplish a defined outcome (ibid.). Computer algorithms operate on information to produce desired outcomes. Nicholas Diakopoulos (2019) explains that there are two primary types of informational tasks of algorithms: intellectual tasks and judgment tasks. Intellectual tasks are tasks with a demonstrably correct answer, such as digitizing document bits into indexable words or computing arithmetic. Judgment tasks are those with no correct answer, in which a preferred alternative is chosen based on facts and attitudes about alternatives. Examples of algorithmic judging decisions include prioritizing, classifying, associating, and filtering. These algorithmic judgments are baked in via rules and procedures articulated by software developers. The primary value proposition of algorithms is the execution of high-quality decisions quickly at scale using automation. By enhancing the scale, personalization, and engagement of content, automated news production is appealing to news outlets in its ability to add to the bottom line of these organizations (ibid.).

It is important here to distinguish between automated journalism and automated news production from a definitional perspective. Automated journalism is defined as computer-written news, with little to no human input beyond the initial programming (Carlson, 2015). Automated news production includes, more broadly, the process in which content is edited and/or moderated by humans before being published. Also, the nuanced distinction between automation and artificial intelligence warrants clarification. Automation and AI are distinct concepts, yet they can work together and complement one another. Automation does not inherently possess

artificial intelligence, but it can if machine learning is in place, allowing for the automation process to take in new data, learn from it and begin to make its own improved and informed decisions (Peng and Bhaskar, 2023). Thus, the term “automated news production” is used in this study to encompass the writing of both template- and generative-based content.

These distinctions are important given the relevance of artificial intelligence tools across the journalistic value chain and the fundamental role of human input in interacting with LLMs like ChatGPT. The journalistic value chain is the entire process of media delivery, from the initial idea to consumption. It includes the various stages and stakeholders involved in delivering media content to the audience. Michael Porter of Harvard Business School noted that creating value is the central activity of successful companies. For any business to survive it must create value for customers more effectively than its competitors (Porter, 2001). The value chain is a framework for analyzing all activities of a firm and then studying the economic implications of those activities. Within a firm, its value chain is composed of five primary activities: inbound logistics, operations, outbound logistics, marketing and sales, and after-sale service (Porter and Kramer, 1985). To apply Porter’s value chain analysis to online news sites, Huang and Heider (2017) identified five equivalent processes: news selecting, news producing, news distributing, news advertising, and interpersonal interaction.

The role that AI plays in each of these processes is multifaceted. A more granular representation of the prevalence of AI across the journalistic value chain is shown below in Table 2.3.1. AI systems are used at the early stages of the value chain to discover information, analyze trends and audience traits, and detect stories. The ideation and research roles of AI correspond to the news selecting process in the online news value chain. The production phase is a mere slice of the pie in terms of value chain, to which the role of AI comprises not just the generation of

content, but formatting, transcribing, summarizing, and translating as well. Both the distribution and advertising elements of the online news value chain incorporate automation in recommending and personalizing content to readers, along with augmentative tasks such as version development and subtitling. After the content is distributed to receivers, news organizations use AI to manage feedback, which is a component of the interpersonal interaction process identified by Huang and Heider (2007), entailing tasks such as comment moderation and audience interaction. Lastly, a critical role of AI, which does not precisely correspond with a particular value chain process, entails automatic tagging and relational analysis.

Value Chain Process	Role of AI	Use of AI systems
News Selecting	Ideation	<ul style="list-style-type: none"> ● Information discovery ● Audience analytics ● Trend/Anomaly/Recurring event detection ● Brainstorming
	Research	<ul style="list-style-type: none"> ● Info validation ● Data set scanning
News Producing	Production	<ul style="list-style-type: none"> ● Content production ● Content formatting ● Transcription

		<ul style="list-style-type: none"> • Summarization • Translation
News Distributing News Advertising	Publication	<ul style="list-style-type: none"> • Develop various versions • Subtitling • Recommendations • Personalization
Interpersonal Interaction	Feedback Management	<ul style="list-style-type: none"> • Comment moderation • Audience interaction
(n/a)	Informational Archiving	<ul style="list-style-type: none"> • Automatic tagging • Relation analysis

Table 2.1: Automated News Production and the Journalistic Value Chain

With regard to bylines for articles produced by automation, different newsrooms have different policies. Computer authorship was identified as a legally complex issue as early as 1965, in which the United States Register of Copyrights distinguished between cases of computers being used as assistant tools to produce articles and cases in which “traditional elements of authorship ... were actually conceived and executed not by man but by a machine” (Montal and Reich, 2017). In the United Kingdom, the Copyright, Designs, and Patents Act was passed in 1988. The act presented a framework for copyright protection of computer-generated works and defined authorship of these works as the “person by whom the arrangements necessary for the creation of the work are undertaken” (Ihalainen, 2018). This definition leaves room for interpretation in terms of different author candidates. In the United States, no legislative

framework has yet to define authorship requirements for computer-generated works, leaving judgment on automated articles to be rendered on a case-by-case basis (Montal and Reich, 2017). As an example of automated authorship indication, when the Associated Press partnered with NLG software company Automated Insights to begin automating quarterly earnings reports using their Wordsmith platform, a note would be included at the end of an article stating “This story was generated by Automated Insights” (Miller, 2015).

As of 2019, “fully automated journalism” was rare and almost strictly used by large news outlets like Reuters and the Associated Press (Karbal, 2019). Automated journalism has historically lent itself best to news stories that rely on the presence of structured data, such as finance reports and sports recaps (Graefe, 2016). Topics like finance and sports can be engaged with in a deterministic, quantitative manner, and most early examples of automated journalism are those in which tabular data from quarterly earnings reports or match results are filled into pre-written templates for a given topic. In quarterly earnings reports published by the Associated Press, for instance, the data upon which the algorithm operates is structured quantitative data, and the algorithm simply has to insert the appropriate data into a pre-existing template for a particular company (AP News, 2021). From a journalistic perspective, this process enables large-scale production of large amounts of similar articles.

In contrast, with articles that require more originality with regard to narrative, human input is more often needed in order to correct mistakes and provide additional contextual input to supplement the A.I.-produced content. As alluded to earlier, in addition to the automated writing of news, automation has been integrated into the news production process in other ways, including data scraping (extracting data from human written output), issuing alerts to trending stories (automatically classifying stories as newsworthy and sending device notifications for

those stories), aggregation and funneling of content (accumulating stories for a personalized audience), data visualization, and the auto-publication of news (automating processes for publishing articles without a human gatekeeper). In all of these processes, technology plays the role of “communicator,” a message source that can create and shape meaning (Wu et al., 2019).

Since the release of ChatGPT in late 2022, a variety of news organizations have led the way in experimenting with generative AI in their news production processes. However, The most publicly salient example of this was Sports Illustrated’s use of AI in story generation, photos, and authors (Bauder, 2023). The media company used stories for product reviews attributed to authors that could not be identified. A picture of one of the authors listed was found on a website that sells AI-generated portraits (see Fig. 2.1). Sports Illustrated said the articles in question were created by a third-party company, and they would be firing the company as a result, but denied that the content of the articles were AI-generated. As a result of this debacle, and reporting surrounding it, Sports Illustrated suffered significant reputational damage (ibid.).

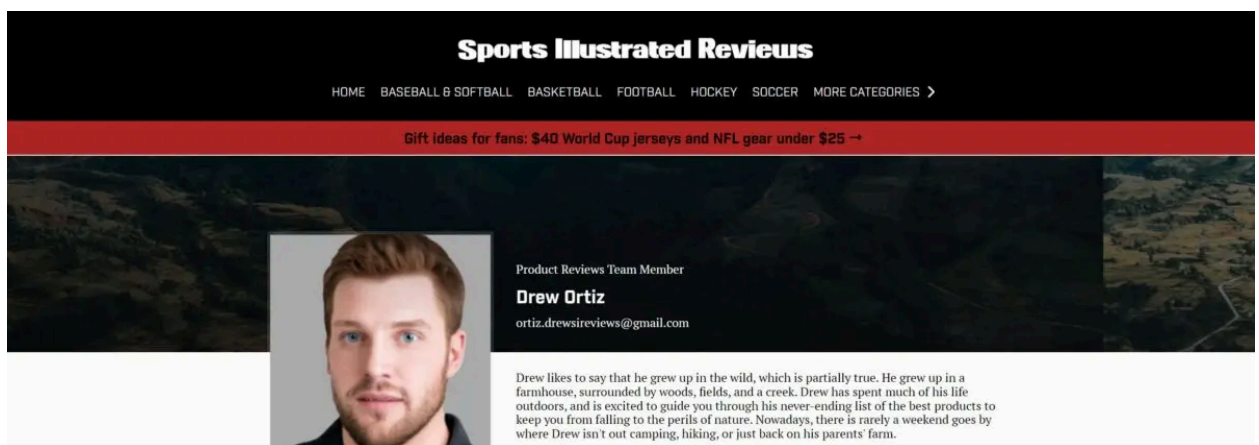


Figure 2.1: An AI-generated author profile at Sports Illustrated (has since been deleted)

CNET and Gannett also experienced backlash after their experiments with generative AI technology backfired. Gannett, a mass media holding company, paused an experiment at some of its newspapers in summer 2023 in which artificial intelligence was used to generate articles on high school sports events, after errors were discovered (see Fig. 2.2). In January 2023, it was reported that popular tech news outlet CNET had used AI to create explanatory news articles about financial topics; these articles were attributed to “CNET Money Staff.” The only way for readers to learn that artificial intelligence was involved in the writing was to click on that author attribution (Bauder, 2023).

HIGH-SCHOOL

Worthington Christian earns narrow win over Westerville North in Ohio high school boys soccer action

LedeAI

Published 10:51 p.m. ET Aug. 19, 2023

The Worthington Christian `[[WINNING_TEAM_MASCOT]]` defeated the Westerville North `[[LOSING_TEAM_MASCOT]]` 2-1 in an Ohio boys soccer game on Saturday.

Worthington Christian edged Westerville North 2-1 in a close encounter of the athletic kind for an Ohio boys soccer victory on Aug. 19.

Worthington Christian drew first blood by forging a 2-1 margin over Westerville North after the first half.

The scoreboard was in hibernation in the final half, with neither team scoring.

You're reading a news brief powered by ScoreStream, the world leader in fan-driven sports results and conversation. Read the [latest exclusive high school sports coverage](#) from the award-winning Columbus Dispatch sports journalists.

Figure 2.2: An example of erroneous AI output in a Gannett newspaper article

The ability of generative AI models to produce freeform text from a corpus of training data has raised legal concerns over copyright issues. For instance, the 2023 New York Times lawsuit against OpenAI highlights the complex landscape surrounding the legality and fairness of AI training practices (Wiggers, 2023). The case, in which the New York Times sued OpenAI and Microsoft for the unpermitted use of Times articles to train GPT large language models, could have a significant impact on the extent to which copyrighted works are permitted to be used in LLM training data sets.

2.3 Why is Journalism being Automated?

Digital technologies in the “era of big data” have contributed to the circulation of an unprecedented amount of information in the realm of journalism (Wu et. al., 2019). Pierre Bourdieu’s field theory is of particular salience when exploring this topic. Bourdieu offers a theory of cultural production in part based on the concept of field, a structured space of positions in which the positions and their interrelations are determined by the distribution of different kinds of resources or “capital” (Bourdieu, 1993). Bourdieu’s field theory identifies key fields within a particular social space and theorizes interconnections and the degree to which they are autonomous of one another. These major fields are the economic and political fields, and their composite field aka the “field of power,” the educational field, the intellectual field, and various cultural fields--journalism being one--which combine to form the “field of production”. The field of power and field of production differ in their levels of economic and cultural capital: the field of power possesses high levels of economic capital but low levels of cultural capital, and the field of production possesses low levels of economic capital and high levels of cultural capital (Hesmondhalgh, 2006). Field theory can produce a comprehensive understanding of the

transformations and struggles taking place within the journalistic field in the age of big data, pointing out that larger political-economic forces are influencing the uptake of automation in newsrooms and causing journalists to react to this trend in distinctive ways (Wu et al., 2019).

Within the field of production, Bourdieu (1993) again makes a distinction, this time between small-scale “restricted” production and large-scale “mass” production. Small-scale production consists of “pure” artistic products. This subfield possesses a high degree of autonomy, but never full autonomy. Furthermore, small-scale artistic producers benefit from very high levels of field-specific symbolic capital. However, they possess very low levels of economic capital, as it is difficult to monetize creative products at a smaller scale. Mass production is, rather than being autonomous, “heteronomous”--meaning it is subject to outside rule. This subfield centers on the production of “commercial” cultural goods, and while it possesses high levels of economic capital through sales, it possesses low levels of symbolic profit and capital (Hesmondhalgh, 2006). The idea of autonomy versus heteronomy is crucial in understanding field-level structural transformations: field autonomy is valuable because it allows for the proper creative process inherent to the particular field and ultimately augments resistance to the “symbolic violence” exerted by the dominant system of hierarchization (Benson, 1999). The shift toward integrating automation into journalism represents a broader trend of the field toward mass production. Optimists about the future of journalism contend that automation will free up time and other resources for journalists to focus on more creative, human interest work. However, within the domain of mass production, it is important to examine the values of the news outlets publishing automated articles and the tech companies that provide these news outlets with content production algorithms. Indeed, with such large-scale transformations of the journalistic field, focus has been increasingly placed on the technological firms that enter into the field and

supply newsrooms with automated technologies. Researchers have also begun to focus on how computer scientists and journalists work together and influence each other. These groups, however, do not operate on the same values: designers of news applications tend to focus on code, personalization algorithms, and relational databases, rather than abide by journalistic standards (Wu et al., 2019).

There are a multitude of external forces incentivizing the trend toward automation in the realm of journalism. External forces including economic recessions, the rise of the Internet and shifting consumption habits, commercialization and consolidation of media, and loss of public trust in mainstream news sources have all weakened the autonomy of the field--in other words, the cultural capital that makes the field unique from other fields (ibid.). The Great Recession left the journalistic landscape with laid-off reporters and the demise of once-great newspapers (McChesney and Nichols, 2011). The rise of the Internet has led to heavy competition for the attention of journalism's target audience through the proliferation of blogs, clickbait, and easy access to information (Jeon and Nasr, 2016). With regard to consolidation, by 1999 six conglomerates controlled most of the media market within the US, and even globally. There are significant consequences of consolidation on the allocation of resources within media firms. While some believe that the deeper pockets of media conglomerates allow them to commit more resources to production, news gathering and so forth, a counterargument is that resources are being diverted from these goals to meet the high cost of financing mergers and acquisitions (Cooper, 2007). At the same time, public trust toward institutions, particularly the mainstream media, has steadily declined. According to Pew, while 76% of the public had either "a great deal" or "a fair amount" of confidence in scientists to act in the common good, only 38% shared that sentiment about the mainstream media (Verma et al., 2017).

Media companies are using AI software to manifest a leaner, more efficient news production process. Shoshana Zuboff of Harvard Business School has argued that “everything that can be automated will be automated” (1985, p. 6). In this sense, automation is part of a broader aim to replace human effort and skill with a single technology that enables the same processes to be performed more cost-effectively and with more control and continuity (ibid.). With regard to automated journalism, everything will be automated if and when it becomes acceptable to the relevant influential social groups. Internally, these include publishers, news managers, developers, and business analysts in media companies. Externally, these include advertisers, the general audience, producers of data, and middlemen between data producers and users (Linden, 2017). The automation of journalism is heavily influenced by what Bordieu calls the “rules of the game,” and the conflict between the heteronomous (mass-production, consolidated media corporations) and autonomous (restricted production) poles of journalism.

Noticeably absent from these groups are news writers themselves. Historically, automation and technological change have increased employment in certain circumstances: for example, when ATMs were developed in the 1970s, the number of bank tellers per branch was reduced, but so was the cost of operating a bank branch, which led to an increased number of branches and thus an increased demand for tellers, who also were able to become more involved in relationship banking. This illustrates an important point about the automation-labor dynamic: workers benefit through continued employment when they can supply tasks that are complemented by automation, but if they cannot provide complementary tasks, they will be at an increased risk of being laid off in a newsroom environment that has integrated automation into its news production process (Autor, 2015). In the era of automated journalism, the increased use of algorithms has placed pressure on journalists to acquire a more technical skill set (such as by

learning to code) in order to increase their economic capital as a means of adaptation to the need for complementary tasks.

Concerns about the extent to which automation substitutes for journalistic labor has manifested itself in the form of what Daniel Akst (2013) calls “automation anxiety”. Akst explains that the proliferation of automation has led to fears about potential impacts on the job market, such as the elimination of jobs and increasing unemployment. Scholars disagree about the threat posed by automation to the field of journalism: for instance, Northwestern University professor Nicholas Diakopoulos contends that while the rise of algorithms in the realm of journalism will change the nature of the work relationship between human journalists and automated systems, algorithms will not replace journalists, but instead would free up higher-level work for journalists by automating more routine tasks. Diakopoulos (2019) estimates that only approximately 15% of reporters’ time and approximately 9% of editors’ time is automatable using currently demonstrated technology, compared to approximately 69% for paralegals, for example. He asserts that the future of algorithms in the domain of journalism will emphasize the design of efficient and effective human-computer systems that free up valuable time for staff to focus on thematic stories, deeper analyses of automated content, and generally more exciting content. Other pundits have noted that computers have no will, no originality, and cannot claim authorship, which nullifies their ability to “replace us” (Cahn, 2020).

2.4 ChatGPT and Natural Language Generation

ChatGPT is a chatbot app powered by the GPT family of AI models that include large language models (LLMs). The app uses GPT models to interact with humans in a conversational way. GPT-4 is the latest Generative Pre-Trained Transformer model created by OpenAI, a San

Francisco-based artificial intelligence research lab. OpenAI was founded as a non-profit organization in 2015 via a \$1B collective pledge from a group of Silicon Valley titans that included Elon Musk and Peter Thiel. Its mission is to ensure that artificial general intelligence (AGI) benefits all of humanity, and pursuant to that mission, its goal is to be the first to develop AGI (Dale, 2020). Artificial general intelligence refers to the ability of an intelligent agent to understand the world or learn a task as well as a human being can (Goertzel, 2014).

The transformer deep learning model, of which OpenAI's GPT models are examples, was introduced in 2017 (Vaswani et al., 2017). Prior to the introduction of transformer models, most natural language processing models were trained on a particular task using supervised learning. Supervised and unsupervised learning are two types of machine learning tasks that differ in how they infer patterns in their training data-- how they, quite literally, "learn". Supervised machine learning requires algorithms to be trained with labeled inputs paired with desired outputs. For instance, a shape with three sides would be labeled as a triangle. The goal of this approach is for the algorithm to be able to correctly predict labels for new inputs that it takes in. Unsupervised learning, in contrast, does not require these labels. The algorithm sifts through unlabeled input data and identifies patterns that will allow it to group the data. For instance, an algorithm may parse through various shapes and group them based on the number of sides they have. In this sense, the mode of learning is unsupervised because there are no teachers and no right or wrong outputs: the algorithm is responsible for analyzing the underlying structure of the data (Sathya and Abraham, 2013).

Supervised machine learning models have two major limitations: they need large amounts of annotated data in order to learn a particular task--this amount of data is often not easily available--and they fail to generalize for tasks other than what they have been trained for

(Kotsiantis et al., 2007). Researchers at OpenAI created GPT-1 in 2018 to address these limitations by proposing a language model that would learn using unlabeled data (unsupervised) but would be fine-tuned using specific examples (supervised). In this case, unsupervised learning served as the pre-training objective for supervised fine-tuning, hence the name “Generative Pre-Training” (Radford et al., 2018). Also in 2018, researchers at Google created BERT, another Transformer-based pre-training model. Another pre-trained model, XLNet, was introduced around this time as well by researchers at Carnegie Mellon. The creators of XLNet directly compared it to BERT and noted its improvements upon BERT on 20 separate tasks (Yang et al., 2019). GPT-1, BERT, and XLNet demonstrated the power of generative pre-training and opened up avenues for other models to realize this potential better with larger datasets and more parameters.

In February 2019, OpenAI announced GPT-2 (Generative Pre-trained Transformer 2), an unsupervised transformer language model larger than its predecessor, with 1.5 billion parameters trained on 40GB of text from the Internet. When used to repeatedly predict the next word in a text based on the preceding context, the model was capable of generating coherent output and making plausible claims, although it was also capable of outputting gibberish (Dale, 2020). GPT-2 earned significant attention in the press as the media picked up on these main aspects of GPT-2: its role as a machine learning algorithm, the plausibility of its output, and its potential for wide societal application. Some mainstream outlets demonstrated the technology by allowing it to be a co-contributor: for example, *The New Yorker*'s John Seabrook (2019) discussed predictive text technology more generally in an interactive piece that, at various points in the article, allows the reader to view GPT-2's contributions based on the preceding human-authored content. *The Economist* used GPT-2 to provide a response to a youth essay question about

climate change with a team of six human judges to assess the results. Given the choice of whether or not to advance the GPT-2 essay response to the next round, four judges responded with “No” and the other two responded with “Maybe” (The Economist, 2019). This result reflects the idea that while GPT-2 technology had the potential to generate eye-popping headlines, its execution still left much to be desired.

Soon following the release of GPT-2, in March 2019, the previously non-profit OpenAI restructured as a “capped-profit” company--a company that cuts returns from investments past a certain point--with the stated reason of attracting more capital. The company would limit profits emerging from investments in the OpenAI limited partnership (the new capped-profit entity) in excess of 100x returns. Open AI claimed that it was unlikely to raise the money necessary to achieve its goals while operating as a non-profit (Coldewey, 2019). The pressure of competing with corporations like Google (creator of the BERT language processing model) in order to remain at the cutting edge of AI research likely made the money issue of particular salience to OpenAI and its leadership. They justified their apparently large 100x cap on returns by referencing their lofty goal of artificial general intelligence and its potential to, if successful, generate orders of magnitude more value than they would owe to their investors. Profits in excess of the 100x cap would go to OpenAI Inc (the nonprofit body that governs the capped-profit entity) (ibid.). In July 2019, following this shift, Microsoft made a significant agreement to invest \$1 billion dollars in OpenAI over the next decade. Microsoft CEO Satya Nadella stated that the partnership would keep “AI safety front and center” so that “everyone can benefit,” and OpenAI CEO Sam Altman stressed that safety and spreading the “economic benefits” were key concerns (Vincent, 2019). The choice of Microsoft as OpenAI’s preferred

partner for commercializing their pre-AGI technologies marked the beginning of a financial collaboration that strengthened following the release of the successor to GPT-2, GPT-3.

In June 2020, OpenAI announced the release of GPT-3. The new language model is more than 100 times larger than GPT-2, with 175 billion parameters and 96 layers trained on a corpus of 499 billion tokens of web content, making it by far the largest language model that has been constructed to date (Dale, 2020). A “token” of web content is a small amount of text, such as a word, part of a word, or just characters; the average token size is approximately four characters (Li, 2020). The number of parameters denotes the complexity of the neural network model. Neural networks are a type of mathematical graph, and in graphs, there are nodes (artificial neurons) and edges between those nodes. These edges can also be weighted. Some neurons are more strongly linked than others, and if the vertices of the graph represent the individual neurons, and edges represent connections between pairs of neurons, then the “weight” of an edge would measure the strength of the connection between two associated neurons (Muller et al., 2012). The number of parameters corresponds roughly to the number of weighted edges in a neural network model. Layers can be understood as a container of neurons that groups a number of neurons together. Neural networks are typically organized in layers: an input layer, an output layer, and any given nonnegative number of hidden layers. Each layer has a particular purpose, and the number of layers in a neural network model corresponds to the complexity of the model: the more layers in the neural network, the more complex the model is (ibid.).

In November 2022, OpenAI released ChatGPT, and the chatbot quickly went viral on social media as users experimented with its capabilities. ChatGPT was fine-tuned on a model from the GPT-3.5 series, a set of new versions of GPT-3 introduced at various times in 2022. ChatGPT has been credited with creating an AI boom that catalyzed widespread public attention

investment toward artificial intelligence (Weise, 2023). It also propelled a large amount of investment into the field of artificial intelligence from venture capital firms and outside investors; as of February 2024, OpenAI was valued at over \$80 billion (Metz, 2024). The chatbot was trained for conversation applications using a combination of supervised learning and reinforcement learning with human feedback (RLHF). In these instances of reinforcement learning, human feedback is commonly collected by prompting humans to rank instances of the agent's behavior. In March 2023, OpenAI launched GPT-4 as part of a paid version of ChatGPT accessible via subscription. OpenAI declined to reveal the exact size of the model but the number of parameters is estimated to be somewhere in the trillions (Kounaa, 2023).

Human-computer interaction researcher Jakob Nielsen (2023) noted that ChatGPT and other generative AI systems have launched the third user-interface paradigm in the history of computing and “the first new interaction model in more than 60 years.” The first UI paradigm was batch processing, where users submitted instructions for the computer to process with no real-time interaction. The second paradigm, command-based interaction, emerged in the 1960s, allowing users to interact with computers via commands, a model that still dominates today through graphical user interfaces (GUIs). Now, emerging with generative AI systems like ChatGPT, the third paradigm is "intent-based outcome specification," where users specify desired outcomes rather than issuing individual commands (ibid.). While intent-based specification significantly reduces the need for detailed instructions, this paradigmatic evolution raises concerns about usability and the loss of user control. The future of UIs may involve a hybrid of command-based and intent-based approaches (ibid.).

2.5 Algorithms: Representation versus Correlation

The increasing application of algorithms in politically consequential contexts has catalyzed efforts to understand emerging digital technologies as tools that produce social ordering in specific ways. The term “algorithmic governance” has been used to encompass a variety of fields and research inquiries relating to the role of algorithms in the carrying out of governmental decisions (Katzenbach and Ulbricht, 2019). Beyond the algorithmic production of news content, examples of such practices include predictive policing, regulatory oversight, and benefits administration (Coglianese and Lehr, 2019). In order to make claims about the political effects of algorithmic media, it is imperative to investigate algorithms as technical objects and the relations they generate. In particular, the nature of algorithmic media as virtual presentation rather than mimetic representation contrasts with traditional forms of media such as writing, painting, and film. The introduction of a virtual ontology into the realm of journalism warrants an exploratory approach beyond an epistemic examination of the normative effects of algorithmic media.

Algorithmic decision-making is inherently different from human decision-making: algorithms, unlike humans, do not (and cannot) give reasons for their decisions (Panagia, 2021). The automated production of information in the realm of journalism therefore involves decisions of a different nature than traditional human-directed news production. Traditionally, mimetic media such as writing, painting, and film, has been understood as representational in nature. These forms of media present reality as we consciously perceive it to ourselves. For example, film transcribes reality in the sense that a physical world that is not immediately present is available for a subject to view. Algorithms, on the other hand, are not technologies that are designed to represent the physical world in this matter: they are correlational in the sense that they model reality through the discrepancy between their generated outputs and a “predefined,

optimal actuality” (ibid.). The algorithmic transcoding of information, in contrast to mimetic representation, need not have a physical correlate in the empirical world. Algorithms are trained on real-world data inputs, but their outputs are simply approximations based on the world of patterns in which they operate.

Thus, it would be unfit to apply the same standards of political analysis to algorithmic media than we would to traditional, mimetic media. Algorithms as political entities warrant an analysis independent of legacy frameworks that have historically been developed and applied specifically toward mimetic forms of media. Cybernetic theory, for example, can provide insights into how algorithms arrange entities in time and space. The distinction between signal and noise is of particular relevance here. Algorithms act as “governors” in the sense that they exert control over the variability in the system (noise) between output and actuality. In the context of automated journalism, algorithms exert control over the words that are deemed relevant to a particular placement in a sentence based on its corpus. What that actuality is, to which the output is calibrated, depends on the values and priorities of the people and/or organizations who designed and developed the algorithm in question. Algorithms are not independent of human choice or shaping; they are designed with some purpose, goal, or outcome in mind.

In September 2020, a few months after the release of GPT-3, Microsoft acquired an exclusive license for GPT-3, which led to criticism regarding the partnership in relation to OpenAI’s goals and overall mission. Elon Musk, a co-founder of the company, criticized the move, tweeting “This does seem like the opposite of open. OpenAI is essentially captured by Microsoft” (Musk, 2020). Other critics noted that the collaboration between the two companies seemed to contradict OpenAI’s aim of making its work accessible to all (Goodwin, 2020). The

consequences of this deal are unclear, but it's likely that the API access will be unaffected, whereas Microsoft's customers might eventually see the benefits of GPT-3 in a range of applications effectively for free (Dale, 2020). Such a partnership indeed begs the question, if OpenAI's mission is to ensure that the benefits of AI are "as broadly and evenly distributed as possible," as to how its decision to exclusively license its state-of-the-art technology to benefit one of the richest companies in the world aligns with that goal.

The issue of OpenAI and Microsoft entering into an exclusive licensing partnership belies a broader issue regarding the opacity of algorithms as a barrier to research and understanding among the general public. In the professional settings in which algorithms are deployed, such as technology and finance companies, algorithms are portrayed as trade secrets--secretive, proprietary items that must be veiled from public view. This phenomenon is not novel: large organizations such as private sector firms and public institutions have long had internal procedures that were not fully understood by the general public to which these procedures were subjected (Burrell, 2016). Yet with new applications for algorithms, particularly in areas that involve public participation and discourse such as social media, an increase in calls for regulation of algorithms and their proprietors has shed light on the opacity of these technologies. The protection provided by trade-secret opacity allows companies like Microsoft to direct attention toward their products and away from the mechanizations that produced them, which comes at the expense of public awareness.

Another manifestation of algorithmic opacity is in the technical literacy required to design and code algorithms. Algorithms are presented as technical skills that require specialized training: for instance, most professions that deal with the creation or management of algorithms require some form of certified knowledge in computer science, data science, or a related field.

The power and authority expressed by algorithms draw from their function as objects in an environment dominated by people with specialized, highly technical skill sets (ibid.). Not only does algorithmic inscrutability preclude understanding by the general public, it also presents difficulties for researchers seeking to understand the social and cultural impact of the implementation of algorithmic systems. Frank Pasquale (2015) writes that while powerful financial institutions and government agencies are able to hide their actions behind non-disclosure agreements and proprietary methods, everything that individuals do online is recorded: “the only questions left are to whom the data will be available, and for how long”. He contends that the implications of concealing information in a black box -- “a system whose workings are mysterious; we can observe its inputs and outputs, but we cannot tell how one becomes the other” -- undermine the openness of our society and the fairness of our markets (ibid.). As a result of this lack of transparency of the algorithms and data collection practices of Wall Street and Silicon Valley institutions, researchers from outside perspectives face an uphill battle in efforts to audit or otherwise render meaningful disclosure from these proprietary tools.

Individuals consume news and other content on social media, but these consumption behaviors are based on information that is often catered to their existing interests by recommendation algorithms. Recommendation algorithms that present content tailored to particular user preferences have been heavily criticized for filtering the information observed by users, who may be placed into biased filter bubbles where the only content they access is the type of content they like and is generated by other people with similar opinions (Pariser, 2011). This comes as a consequence of the fact that recommendation algorithms emulate a feedback loop in which the system aims to reinforce a cycle that attempts to optimize user retention and interaction. These algorithms attend to user engagement signals such as topic preferences, social

connections, and topic relatedness, to guide the recommendation of information, and are therefore affected by popularity and homogeneity biases (Chaney et al., 2018). In this context, the biased filter bubbles within which users consume information may not only limit their exposure to diverse points of view, but also reduce the quality of the information they access, potentially making them vulnerable to the proliferation of false information (Spohr, 2017). In particular, social media platforms have been criticized for amplifying content that is divisive, sensational and conspiratorial (Tucker et al., 2018).

2.6 Previous Studies on Automated News: Reach and Practice

Main finding #1: A 2021 systematic review of automated news scholarship by Samuel Danzon-Chambaud found that inquiry used in automated journalism scholarship could be categorized under two main fields: the reach of automated journalism and the practice of automated journalism

In a 2021 systematic literature review, Samuel Danzon-Chambaud of Dublin City University analyzed the key features of a selection of academic articles on automated journalism in order to provide a comprehensive overview of the field and contribute guidelines for future research. The intent of the review was to focus on the full range of peer-reviewed journal articles contributing to the growing research area of automated journalism beyond just empirical studies on reader perceptions of automated news (Danzon-Chambaud, 2021). Within the academic community, Danzon-Chambaud suggests, automated journalism as a technology is sometimes discussed along with the algorithmic distribution of media content. The technology can also be

investigated as part of computational journalism studies. Thirdly, automated journalism can be studied in the context of newsroom automation and artificial intelligence in journalism, which looks at media industry breakthroughs and platform distribution.

Based on a keyword analysis, Danzon-Chambaud makes an important distinction between the term “automated journalism” and the term “robot journalism”. He observes that “robot journalism” intimates a conceptualization of a robot typing on a computer keyboard that ultimately plays on journalist fears of being replaced and prevents newsroom innovation. He also notes that the term “automated journalism” should not be narrowly conceived: a definition of automated journalism that focuses on computer-generated text is too narrow and fails to reflect the views of media practitioners. Danzon-Chambaud stresses the importance of encompassing other algorithmic tasks in journalism beyond text generation such as retrieving newsworthy data in investigative reporting or automated fact-checking (ibid.).

More broadly, the term “automation” is one that is notoriously difficult to define, as it possesses a multitude of histories and cultures that have influenced its usage over time. Even among experts, there is difficulty in defining the term in their language, or their region of the world, or their professional domain, as automation has a unique meaning that may not be the same meaning for other experts. Furthermore, it is imperative, particularly within the practice of automated news, to delineate between journalistic practice and news practice. Journalistic practice is not the same as news practice: the practice of journalism entails the activity or profession of being a journalist (writing, editing, etc.), whereas the practice of news centers on reports of current events and other new information of interest broadcast via various forms of media.

Danzon-Chambaud found that inquiry used in automated journalism scholarship could be categorized under two main fields: the reach of automated journalism and the practice of automated journalism. The reach of automated journalism includes studies on the perceptions of news readers and their comparisons to human-written content as well as studies on the implications of automated journalism for domains such as the legal and financial realms. The practice of automated journalism consists of technically oriented studies that examine the functioning of automated text production, studies examining the deployment of automated journalism within news organizations, and studies focusing on media labor and the perspectives of media practitioners. Technically oriented studies encompassed under the “practice of automated journalism” field primarily examine the functionality of the technology to demonstrate both its potential as well as its limitations. Media labor and newsroom-oriented studies in this category on issues such as human intervention in the editorial process, chatbots and media audiences, and attribution bylines.

Among the 33 scholarly articles examined in Danzon-Chambaud’s meta-analysis, the choice of methods often differed between the reach studies and the practice studies. Methods employed for studies on the reach of automated journalism encompassed mostly experiments along with some content analysis approaches. For the studies that examined the practice of automated journalism, the employed methods included content analysis, surveys, interviews, participant observation, and several mixed-methods approaches. The results of the studies on reach show that, with regard to reader perceptions, automated and human-written news content are perceived similarly in terms of the metrics of credibility, objectivity, and trustworthiness, but not as far as reading for pleasure is concerned. Danzon-Chambaud recommends that future

studies examine the “latest breakthroughs in NLG production” against the criterion of reading for pleasure.

Main finding #2: A 2020 meta-analysis organized existing experimental survey-based studies on automated news perception based on the following aspects of the study: participants, stimulus (articles), experimental design, and outcome variables

A predominant subset of existing studies on perceptions of automated journalism are experimental studies in which news content was presented as a stimulus to Internet users measuring some or all of the following outcome variables: credibility, quality, and readability. In a 2020 meta-analysis of 11 studies, all published between 2017 and 2020, falling under such criteria, Graefe and Bohlken observed no difference in readers’ perceptions of credibility, a small advantage for human-written news in terms of quality, and a large advantage for human-written news in terms of readability overall across the 11 studies (Graefe and Bohlken, 2020). These studies presented recipients with a short news story, in which either the author (journalist or algorithm), the attributed author (journalist or algorithm), or both were experimentally manipulated. Recipients would then rate the article they had just read in terms of credibility, quality, and/or readability. Sample sizes for each experiment ranged from 100 to 1000 and were typically restricted to a single country, as recruiting participants within a national context is practical in the sense that it allows for the experimental stimuli to be presented in a single language. Common topics for articles presented to subjects included sports, politics, finance, and breaking news (e.g. earthquake reports). The experimental designs varied between single- and

multi-factor analyses; such factors include author (who actually wrote the article), attribution (the source to which the article was attributed in the stimulus), topics (e.g. sports, finance, etc.), and media outlets (traditional or online).

Main finding #3: Galton and Ruge's 1965 study was important in its identification of news values. Following studies have built on its theoretical foundation, which demonstrates the importance of consistently evaluating the value of news as it evolves.

The criteria traditionally used to evaluate news value, derived primarily from research on human-written content, may require reconsideration. In their landmark study, Johan Galtung and Mari Ruge (1965) significantly advanced the structural understanding of news media. Their central question, “How do ‘events’ become ‘news’?”, led them to identify twelve key factors that play a pivotal role in this transformation. Their research focused on how overseas events were reported as foreign news in the Norwegian press. These factors, ranging from frequency and threshold to negativity and reference to elite nations and people, provided a framework for understanding the selection and distortion processes in news reporting.

Galtung and Ruge's study was pioneering in its approach, suggesting that certain characteristics made events more likely to be selected as news. These factors included the intensity and clarity of an event, its cultural proximity to the audience, and its unexpectedness or continuity in the news cycle. Their hypotheses – that the more an event aligns with these factors, the more likely it is to be reported, and that these characteristics are accentuated in the news-making process – were insightful in deciphering the complex dynamics of news reporting. Galtung and Ruge's work, while hypothesizing rather than demonstrating these factors, laid the

groundwork for future research, encouraging a critical examination of how news is crafted and urging journalists to counteract these influences for a more balanced and comprehensive reporting.

In the decades following Galtung and Ruge's seminal study, the academic discourse surrounding news selection and presentation has continued to evolve, drawing both support and critique from various scholars. Joye et al. (2016) provide a comprehensive overview of this ongoing scholarly conversation. They note that while some studies have confirmed Galtung and Ruge's findings, others have raised methodological concerns and proposed additional news factors, advocating for a more nuanced and expanded model of news selection. The impact of digital media, a factor not present in the original study, has significantly altered journalistic practices, necessitating a reevaluation of traditional news values. Joye and colleagues highlight three critical areas for contemporary news value research: assessing the relevance of Galtung and Ruge's hypotheses in the context of today's data-rich environment, integrating the changing societal and cultural contexts in news selection, production, and reception, and aligning the study of news values with the realities of global journalism. This forward-looking agenda underscores the need to adapt and expand the original model to suit the complexities of modern news reporting, taking into account the proliferation of digital media and the global interconnectivity that shapes current journalistic practices.

Davide Panagia (2023) describes technologies such as ChatGPT as "systems that barter in Bayesian probabilities rather than mimetic representations" -- and thus, the frameworks we use to critique or critically think about these technologies are obsolete insofar as "we think of the activity of 'challenging' as a critical operation that negates a representation or an identity". An approach to critically evaluating media that accounts for this distinction would thus benefit from

an exploratory understanding of how the intended audience of such media perceives it. If news media is no longer a representation of mind, but rather an accumulation of probabilistic calculations, readers are dealing with a different entity of media that is nonetheless presented as possessing those same intrinsic qualities. Moreover, if readers are cognizant of the AI authorship of content, their own evaluative criteria might shift, necessitating the development of new metrics within academic research to aptly compare human and AI-generated news.

For example, the news values identified by Galtung and Ruge may be less relevant in terms of their prioritization by news consumers in the context of AI-generated content. This difference in importance to consumers is independent of how they are traditionally prioritized for marketing and administrative purposes in the newsroom. Existing research on reader perceptions has relied on Likert-type or semantic differential scales that use adjectives deemed relevant by researchers and force receivers to rate news articles and sources along the dimensions researchers propose. However, the relevant psychological dimension(s) along which participants vary in response to stimuli may be different than the adjectives deemed relevant by researchers given the potential differences between AI-generated content and human-generated content in terms of the factors to which readers attend. These dimensions could shift for two reasons: 1) if the content or message itself is different, but also 2) if the content is explicitly attributed to a different source.

Main finding #4: Qualitative studies on automated journalism outside of the experimental survey framework have focused on the perspectives and experiences of study participants. Likert-scale questions can be constraining in the scope of responses they elicit.

Outside of the quantitative experimental framework utilized by many empirical studies on the reach of automated journalism, qualitative research on the practice of automated news production has focused primarily on the journalistic perspective and understanding the impact of automated journalism on journalistic practices and experiences. Dorr (2015) conducted semi-structured interviews in a study that aimed to uncover the economic potential of NLG in journalism and interrogate its institutionalization on an organizational level. Thurman and Dorr (2017) conducted workshops with journalistic practitioners to investigate journalists' judgements on the limitations of automation and its potential to become more common and increase the depth, breadth, specificity, and immediacy of information available. Van Dalen (2012) sought to understand how human journalists define their own skill sets within the broader context of the journalistic profession, and how they see automated journalism shifting those roles. To that end, the author conducted a critical analysis of 68 newspaper articles and blog posts covering journalistic reactions to the launch of a network of machine-written sports websites.

In an investigation of print and online news perceptions in general, Sundar (1999) noted the limitations of quantitative scales in experiments in terms of measuring individual values and perceptions. The study prompted respondents to read a news story and asked them to list the thoughts that came to their mind. Following this, they were asked to provide adjectival descriptions of the content (*ibid.*). Communication research has prescribed a wide variety of attributes or characteristics to use as evaluation criteria for news stories, such as timeliness, potential impact, prominence of people involved, proximity to audience, and novelty of the event (Mencher and Shilton, 1997). Other criteria, such as accuracy (verification of published information), attribution (proper source identification), balance, fairness, objectivity, brevity, and clarity emerged as part of an era of social responsibility of the press (Sundar, 1999). Most of

these criteria are synonymous with credibility, quality, or both. This assumption of social responsibility has guided the integration of these news qualities as dependent variables administered to subjects on quantitative Likert scales ranging from three to ten points. Survey respondents are expected to evaluate and respond to these news stories based on the predetermined variables that are designed to elicit such evaluations.

The presentation of these evaluation measures as quantitative Likert-scale questions could potentially limit the scope of respondents' answers. People may not always carry values in their heads for the particular evaluative measures in which researchers are interested (Chaffee and Schleuder, 1986). Researchers do not definitely know the most relevant dimensions along which individuals vary in their perception of news contents. Qualitative, open-ended questions allow respondents to articulate unprompted variables that describe the self-reported effect of news stories upon themselves. As opposed to quantitative scalar measures, such questions do not force respondents to rate news articles and sources along dimensional variables proposed by researchers. Rather, an open-ended qualitative mode of inquiry gives participants the opportunity to offer an unfiltered indication of the relevant psychological scale(s) along which participants vary in terms of their response to a stimulus (an article, in the context of this study).

In his systematic literature review, Samuel Danzon-Chambaud advocates for the importance of theories of institutionalism as well as Bordieu's field theory as frameworks for studying automated journalism. Many existing studies on the reach of automated journalism apply frameworks derived from the field of psychology, such as selective exposure and expectancy theory to guide their experiments. As for studies on the practice of automated journalism, a common feature among them is the use of sociological frameworks with a focus on media labor (Danzon-Chambaud, 2021). Other authors have pointed toward institutional theory

to explain how a social constructivist approach could help investigate the consumption and production of automated news content (Napoli, 2014).

2.7 Research Questions

Based on previous research and gaps in the literature, this thesis seeks to answer the following research questions:

RQ1: How do individuals involved in automated news production describe their experiences?

RQ2: What are the evaluative criteria used by readers in their perception of automated news content?

RQ3: How do news readers in the United States perceive automated news content and human-written news content relative to the descriptors derived from RQ2?

Chapter 3: Theoretical Considerations

3.1 Overview

After identifying gaps in existing literature and introducing the research questions at the end of the previous chapter, this chapter introduces the theoretical considerations and approaches deemed best suited to approach said research questions of the study, which cover ontological and phenomenological dimensions of automated news content in terms of two dimensions: production and consumption (writing and reading). As such, the necessary historical and theoretical considerations inform both processes.

RQ1: How do individuals involved in automated news production describe their experiences?

RQ2: What are the evaluative criteria used by readers in their perception of automated journalism?

RQ3: How do news readers in the United States perceive automated journalism and human-written journalism relative to the descriptors derived from RQ2?

To foreground each of the three research questions, given that they each address the concept of automated journalism, an ontological assessment of automated journalism is key. This chapter elucidates the importance of understanding the mode of existence of an algorithm in order to understand its effects. In the context of automated news as a system, an ontological assessment involves understanding the “being” of automated systems - what it means for an automated system to exist and operate, and how its existence is defined in relation to humans and other entities. An ontological assessment of algorithmically generated media deals with what it is as a category separate from other categories of media and underscores the importance of

assessing this difference in the context of this study, given the emergence of large language models (LLMs) as sources of written content. Drawing from the concept of the algorithm dispositif, this chapter highlights the underlying ontology of AI-generated textual outputs in particular as a fundamentally distinct entity from the mimetic ontology that grounds contemporary thought and criticism. An algorithmically generated output differs fundamentally from an output written by a human: it is a rendered output rather than a representation of an idea. This is important given the novelty of such a form of media.

In addition to the production of the output and its ontological nature, it is crucial to contextualize readership as an activity in and of itself, which entails the construction of meaning of a text and bears the weight of readerly expectations and responses. The second and third research questions, RQ2 and RQ3, concern the question of readerly reception: one's relationship to writing and reading. Thus, this chapter makes use of reader response theory to detail this relationship in terms of the creation of meaning in a text, and its consequent implications for automated text consumption versus human-generated text. The ontological world of the algorithm changes the expectation of readership and readerly reception that are available to readers due to a shift in the locus of trust brought on by these evolving technologies. As such, this chapter makes use of Nietzsche's theory of the transvaluation of values and Hume's theory of moral sentiments to detail how changes in text generation technology compel our moral sentiments to shift.

As it pertains to the subjective experiences of individuals involved in the production of automated news, the first research question necessitates phenomenological considerations. Similarly, for the second and third research questions, phenomenological insights would also be valuable given the focus on reader perceptions in their assessments of AI- and human-generated

news articles. This chapter therefore emphasizes the importance of extending beyond just an ontological assessment of the form of media itself to more deeply understand human experiences with such media from an epistemological and phenomenological perspective, given the critical nature of human involvement in the production and consumption of such media. Polack's theory of paramediation is evoked here as a means of looking beyond an ontological assessment of the nature of this entity to a phenomenological assessment of how people experience it. Such a phenomenological view of algorithms, in contrast to traditional epistemological assessments, is needed to challenge the idea that the social consequences of algorithms are exhaustively reflected in their logic.

As an illustrative example of the ontological and phenomenological dimensions of AI-generated text, this chapter presents a case study of ChatGPT as a multifaceted tool. ChatGPT represents a distinct type of entity in the world of AI. Such a case study is important in that it reflects ontological inquiries into what constitutes an AI entity and phenomenological inquiries into AI as a mediator of experience between information and the user, shaping how information is received, understood, and utilized. This chapter ultimately frames the rest of the study as an exploration within and across the full chain of AI-generated content: what it is, how it is experienced, and how it is produced.

3.2 The Algorithmic Ontology

An ontological assessment that evaluates the role of algorithms in news production is of paramount importance in order to demonstrate the theoretical stakes of AI-generated text as a novel entity. The dispositional powers of automated systems inform their unique capabilities and limitations in arranging, distributing, and presenting information. Drawing on Panagia's concept

of the “algorithm dispositif,” the forthcoming analysis will explore how these technological mechanisms not only function in producing content but also in shaping the nature of the information itself. Moreover, it is essential to recognize the political effects of language generated by algorithms, which operate in a virtual mode of existence rather than merely representing reality as traditional human-driven forms of media do. The distinction between the virtual and the representational underscores the transformative impact of algorithms on the landscape of news production, influencing both the content and its reception.

Understanding the mode of existence of an algorithm is crucial for comprehending its effects. The concept of modes of existence is central to Gilbert Simondon’s theory of how technical objects are integrated within human culture and society. Simondon’s approach encourages examining algorithms on their own terms, acknowledging their specific functionalities and the internal logic that drives them (Simondon, 1989). Assessing the being of automated systems involves exploring what it means for such systems to exist and operate within our socio-technical landscape. The modes of existence of these technical objects are significant because they generate specific relations and transformations in how we perceive, experience time, and associate with each other and our environment.

The challenge of thinking about the “being” of algorithms extends to thinking politically about algorithms, questioning how an algorithm functions as a medium and, more critically, as a political medium. These inquiries necessitate a broader understanding of media and their intrinsic political dimensions. Just as we might consider a guitar amplifier as a communication tool that shapes sound, algorithms can be seen as instruments of transmission—or influence machines, invoking Victor Tausk (1933)’s terminology—that mold information and influence. Tausk describes the influencing machine as a “delusional instrument” that “serves to persecute

the patient and is operated by its enemies”. Tausk further notes the opaque, black-box nature of the influencing machine and its perception by the patient as hostile or diabolical. While we tend to think about algorithms as influencing machines as Tausk describes them, which stems from dominant thinking in critical disciplines, an exploration of algorithms politically is not simply about recognizing the functional parallels between them and instruments of transmission but also critically assessing whether all media wield political power in analogous ways.

Davide Panagia, Louise Amoore, and Peter Polack have contributed significantly to understanding algorithmically generated media as a distinct category, emphasizing the need to perceive algorithms as automated systems that continuously arrange and rearrange relations. The agency of algorithms is based on their dispositional power to order bodies and energies within space and time, operating in a virtual mode that deals more with probabilities and correlations than with verifiable truths. Algorithms therefore represent a different sort of medium, one that challenges traditional notions of representation and truth.

An ontological assessment of algorithmic media extends to the realm of political judgment in an era dominated by algorithmic governance. Unlike human judgment, which is typically framed within binaries like true/false or right/wrong, algorithms ground decision-making in distributed cognition and operate through probabilities, correlations, and the infinite recursion or recombination of outcomes. Louise Amoore uses an intriguing analogy to describe the operational logic of algorithms: algorithms manage information through arrangements based on proximities, distances, intensities, and associations. Consequently, algorithms do not function to represent truths or falsities about the world but rather to create new forms of order and understanding, which requires a fundamental reevaluation of how we interpret and interact with these systems (Amoore, 2020).

Assessing the distinctive nature of LLMs such as ChatGPT as sources of written content is crucial in the context of this study, particularly because these systems operate on Bayesian probabilities rather than mimetic representations. As a result, these systems demand alternative forms of understanding and reception, and a reevaluation of the modes of critical political science thinking that are applicable to understanding these systems is therefore necessary. LLMs embody a form of politics through the types of power they manifest, which extends beyond traditional frameworks of understanding media influence. The dominant critique of ChatGPT, for instance, is that it poses a moral hazard as it challenges long-standing commitments to individualism and autonomy by potentially guiding user responses and shaping thought patterns. However, Davide Panagia (2023) notes that this often overlooks inherent biases such as Anglo-American exceptionalism and neurotypical biases embedded within these critiques.

Panagia (2020)'s concept of the "algorithm dispositif" is critical here as a theoretical framework with which to understand the intricacies of politicality and power dynamics inherent in automated systems. Originating from Michel Foucault's discussions on Manet's paintings, the term "dispositif," often translated as "apparatus," is used by Foucault to articulate the mechanisms of discipline, security, and governmentality within technical media (Bussolini, 2010). By opting for "dispositif" instead of the more conventional "appareil"--meaning apparatus--Foucault conceptualizes the nexus of media, aesthetics, and politics differently. Foucault's terminology marks a shift in political, aesthetic, and methodological parameters, steering away from viewing media objects merely as tools of domination, and instead understanding them as sentimental instruments that actively arrange dispositions, attentions, and perceptibilities, thereby reshaping our engagement with the world and its governing systems. The framework of understanding algorithms through the algorithm dispositif is particularly useful in

analyzing how algorithms, as modern technical objects, influence and shape our lives and associations.

Thinking of algorithms as “sentimental instruments” rather than merely tools of transmission or domination give us the ability to reveal the deeper ontology that underlies their function. More specifically, algorithms are not just neutral or passive entities but active media that shape our perceptions. By coordinating attention and awareness, algorithms alter our existence in terms of our relations to one another, crafting new forms of social interaction and engagement. These arrangements of the world are not merely organizational but are embedded with power dynamics that may include domination—which is a nuanced but critical difference from the conception of algorithms as tools of domination themselves. Adopting the framework of the algorithm *dispositif* encourages a broader, more nuanced understanding of algorithms, moving beyond the consideration of algorithms in general or any single algorithm.

The political ontology of the algorithm *dispositif* differs from traditional representations found in earlier technical media like writing or painting in two major ways: their virtual mode of existence and the political physics of their vital processes. Such a difference signals a shift in how we understand sources of information, and consequently, power and control within society. First, unlike their predecessors, algorithms operate in a virtual mode of existence. They do not generate outputs that are strictly true or false; instead, they create continuously adjustable thresholds that adapt to given inputs. Insofar as algorithms operate within virtual realities—managing probabilities and correlations instead of verifiable truths—they therefore pose a challenge to traditional political theories that are heavily reliant on representational structures.

Second, algorithmic ontologies can be understood in terms of their political physics of vital processes. They possess the capability to control and organize the movement of energies and resources in space and time through automated mechanisms. This is often managed through negative feedback systems, where algorithms govern the movement of people and resources without direct exertions of domination or disciplinary powers. This reconfiguration of how control is exercised challenges conventional notions of political power, suggesting a new paradigm where control is less about overt domination and more about the subtle guidance of systemic behaviors and interactions. This shift prompts a reevaluation of the underlying forces at play in the digital age, where algorithms are not merely tools but active participants in shaping socio-political landscapes.

A historical analysis of mimesis can help us understand the evolution of language and other technical media in the context of representative technologies, illuminating this evolution through the philosophical perspectives of Plato and Aristotle. Plato argued that art, as a form of imitation, is thrice removed from truth. He posited that the material world itself is an imitation of the perfect world of Forms; hence, art, which imitates the material world, is essentially a copy of a copy. Plato's argument serves as a foundation for the aforementioned contemporary critiques of technical systems as influence machines, suggesting that thinking critically about the political nature of technical experiences demands an ability to reflect and rationalize, turning away from the coercive power operations often embedded within technical objects.

Aristotle's contributions further this exploration by focusing on theater and mimetic representation, emphasizing how mediums operate and the forms of association they introduce. His analysis of the actions within a plot and the broader role of mimetic representation in theater highlighted the importance of understanding the purpose of a medium as well as its mimetic

capabilities. In contrast to the prevailing notion during Classical Greek times that associated poetry primarily with music and song, Aristotle argued for a literary understanding, insisting that poetic criticism be oriented toward understanding the literary purpose of a work or medium. Insights from both Plato and Aristotle regarding the mimetic nature of a medium thus enhance our understanding of how both ancient and modern forms of media create associations and convey meanings, shaping the ways in which they are perceived and interacted with in their respective cultural contexts.

3.3 Journalism and Language

An examination of the history of language illustrates the difference that an ontology of AI-generated textual outputs introduces to the representational nature of human-written language. Human language fundamentally uses symbols—such as words or signs—to represent objects, ideas, and relationships. These symbols are arbitrary; for example, the letters C-A-T arranged in that order do not inherently constitute a cat, but humans use this formation to refer to a specific animal in the world. This symbolic and arbitrary nature of language enables the representation of not only tangible objects and actions but also abstract concepts like freedom, love, or justice. Moreover, the grammatical structure of language dictates how these symbols can be combined, allowing for the creation of an almost infinite variety of expressions from a finite set of elements.

Alternatively, algorithmically generated text does not originate from a naturally evolving system of symbolic representation but is instead produced by algorithms that are processing vast datasets. AI systems generate language based on patterns and probabilities identified in their training data rather than direct representation, introducing a new ontology where the connection

between sign and referent is not governed by historical or cultural semantics but by computational processes. As a challenge to our conventional understanding of representation and meaning in language, generative artificial intelligence provides us with an introduction to a future in which textual outputs may no longer adhere strictly to the traditional rules and associations that have defined human communication for millennia.

Human languages operate based on a shared understanding among their users, where people, often implicitly, agree on the meanings of specific words or signs. This consensus allows for effective communication, though the representational nature of different languages varies widely, influenced by the culture and context of their speakers. Such diversity in linguistic frameworks leads to variations in how concepts are understood and expressed across different societies. For example, certain words or phrases might carry unique cultural significance that is particularly embedded in one language but entirely absent in another. Language not only serves as a tool for communication but also requires active interpretation by its users. The representation of ideas or objects in language can be perceived differently by individuals depending on their personal perspectives and cultural contexts. It is the subjectivity of language that consequently highlights its complexity: the same word or phrase can evoke varied interpretations and reactions among different people. Thus, while language is a fundamental medium for sharing information, its effectiveness and clarity are heavily dependent on the shared knowledge and interpretative skills of its users, evincing the intricate relationship between language, culture, and individual perception.

Traditional journalism—*specifically that which does not employ artificial intelligence in the text writing process*—is deeply entrenched in the system of representation inherent in language. Journalists engage in selective representation, deciding which events in the world are

noteworthy or relevant to their audience, thus shaping public discourse. They frame stories by quoting particular sources and using particular language, techniques that serve to shape audience understanding in specific ways. This framing acts as a form of representation that can emphasize certain aspects of a story while omitting or downplaying others, ultimately influencing how the information is perceived and interpreted. Moreover, journalists construct narratives by carefully ordering facts and events in a manner that is coherent and understandable to their audience. Inherently, this process of narrative construction is dispositional, requiring a structured arrangement of information that not only presents the facts but also fits them into a compelling framework. Narrative construction goes beyond merely demonstrating proof or providing evidence. It involves significant interpretive work that highlights the inherent subjectivity in journalism. Journalists must interpret events and make decisions about which details to emphasize, which perspectives to foreground, and how to frame the stories they tell. This interpretive aspect is influenced by the limitations of language, which can restrict how effectively information is conveyed.

Moreover, subjectivity in this process does not end with the journalist, as it extends to the audience as well. Each member of the audience brings their own perspectives, biases, and background knowledge to their interpretation of the news, affecting how they understand and react to the narratives presented to them. This dual layer of subjectivity, both from the journalist in the crafting of the story and from the audience in its reception, underscores the complex interaction between reporter and reader in the shaping and understanding of news. The writer-reader interaction is of paramount importance to the dynamics of modern journalism, where—as we will analyze more deeply later in this chapter—both the creation and consumption of news involve a continuous negotiation of meaning and perspective.

Mimesis significantly influences traditional journalism, with the practice of human-written content creating representations of the world for readers. Technical advancements such as the advent of online news sites have shifted journalistic practices from face-to-face interactions to mediated monitoring. Journalists now extensively monitor competitors and other news outlets via digital channels, often leading to risk aversion where there is a tendency to mimic stories and approaches that have already proven successful elsewhere. As a result, increased content overlap has proliferated across media outlets, which is especially evident in the stories that dominate front pages and news feeds.

The materiality of contemporary journalistic practices has also been transformed by the growth of online publishing, which necessitates constant content production and has intensified mediated monitoring. The creation of an infrastructure of mediation, which incorporates various hardware, software, and connectivity devices, allows journalists to actively source news from multiple platforms such as wire services, online news sites, and social media. When mediated monitoring processes and imitation practices converge, a materiality-mimicry connection is formed, which is evident in contemporary journalism in terms of how the physical and digital tools employed by journalists shape the news landscape through a cycle of continuous observation and replication (Boczkowski, 2009).

The virtual mode of existence of algorithms stands in contrast to the mimetic reality of human-generated texts. Where human texts represent ideas and can typically provide reasoning behind the choice of words and constructions, algorithmically generated texts merely render outputs without such explanations. Algorithms inherently lack the capacity to give reasons for their decisions, a limitation evident across various fields including finance, predictive policing, and sentencing (Panagia, 2021). This limitation is also true in journalism, where algorithms

cannot provide reasons for their choice of a particular word or sentence. The importance of distinguishing between mimetic and virtual realities is underscored by the novelty of algorithmic forms of media. Unlike journalists who write sentences in news articles that convey fixed meanings, AI-generated text outputs probabilities, indicating a fundamental difference not just in the nature of the output but also in the process of text generation itself.

3.4 Readerly Expectations and Values

RQ2 and RQ3 relate to the dynamics of readerly reception—the intricate relationship individuals have with writing and reading. Traditionally, most readers learn to read through established educational systems, which are deeply rooted in specific pedagogy and traditional ontologies of teaching reading, including the methods of literary criticism. However, the current teaching methodologies do not account for the emergence of automated writing technologies. The mismatch between traditional reading expectations and the capabilities of new technologies like ChatGPT presents a significant challenge, particularly for educators such as professors who find their teaching paradigms tested by these innovations. A key issue is that unlike human authors, automated texts like those generated by AI do not possess agency or responsibility, and therefore, cannot be attributed with intention. The fundamental question therefore is how readers interpret and interact with text generated by AI (and furthermore, how they *should* interpret and interact with this text) as traditional interpretative strategies may not directly apply.

The literary theory of reader response criticism centers on the idea that meaning in a text is actively created by the reader. When a text is consumed, the locus of meaning could be the writer, the text itself, or, most critically, the reader. The focus of this theory is on the reader and their engagement with the text, emphasizing that meaning does not exist independently but

emerges through interaction. The complexity of reader response criticism lies in the diversity of reader perspectives and responses, which vary widely based on individual backgrounds and experiences. Ultimately, variability in these backgrounds and experiences raises the question: how can one craft a meaningful piece of reader response criticism if responses are so individualized?

To navigate this, reader response theorists consider what an ideal reader might derive from the text, what audience the text aims to reach, and how such an audience might react. They explore how the text might be interpreted from various perspectives and how it invites responses from its readers. As one reads, the text generates certain expectations, and part of the critical endeavor involves discussing these anticipations. Are these expectations fulfilled, redirected, or denied as the narrative unfolds? The process is inherently dynamic, and the most compelling writing often surprises us by subverting our expectations. The key takeaway here is that such an interaction is different between readers and AI-generated texts, which do not create the same anticipations and emotional engagements as texts written by humans.

Gertrude Stein's innovative use of automatic writing and Mark McGurl's analysis of literary institutions have both influenced shifts in literary theory towards a greater emphasis on reading practices and the factors that shape them. Stein used automatic writing as part of her broader experimental approach to literature, which sought to redefine the relationship between words and their meanings. Automatic writing is a technique in which a person writes without consciously intending what to write, allowing thoughts to flow onto the page without premeditation or self-censorship. This method is often associated with spiritualism and psychological exploration, where it is used to uncover thoughts and feelings that are not accessible through ordinary conscious thought. In a literary context, such as that in which Stein

experimented, it serves as a means of tapping into the subconscious mind to access raw and unfiltered creativity (Will, 2021).

Stein's work, characterized by its ambiguity and open-endedness, often places greater demands on the reader, requiring active interpretation and engagement. This aspect of her writing ties into the aforementioned shifts in literary theory, particularly reader-response criticism, which emphasizes the role of the reader in creating the meaning of a text. Stein's experimental use of language, with its unconventional grammar, syntax, and form, required readers to approach her texts in a manner that is different from more conventional narratives (Friedman and Fuchs, 1989). This challenges traditional reading practices, pushing readers to develop new strategies for comprehension and engagement.

In "The Program Era," McGurl argues that the expansion of academic creative writing programs has significantly shaped modern American literature. These programs have influenced the styles and themes of literary works by creating a sort of institutional framework within which writers develop and refine their craft (McGurl, 2005). His analysis of how creative writing programs shape literary production indirectly engages with ideas central to reader-response criticism, such as the role of the reader and the interpretive communities that influence how texts are understood and valued. One of the key ideas in reader-response criticism, developed by Stanley Fish (1976), is that of "interpretive communities," which are groups of readers who share particular interpretive strategies. McGurl's examination of the academic and workshop environments in creative writing programs essentially looks at how these settings form their own interpretive communities. These programs teach specific ways of reading and writing, thus shaping the literary preferences and critical approaches of their participants.

While McGurl focuses more on the production of literature within academic settings, the implications of his work suggest that the ways texts are taught and critiqued in these programs influence how readers (students, teachers, literary critics) create meaning from texts. This aligns with reader-response criticism in its emphasis on the active role of the reader in constructing the meaning of a text. The shift from human to AI authorship changes the ways texts are taught and critiqued due to the algorithmic ontology and thus alters how readers create meaning from these texts.

The algorithmic ontology introduces a new element to the landscape of readership and readerly reception, shifting expectations from engaging with representations to navigating automated inferences. Inferences are now not directly crafted by human intent but instead are algorithmically generated based on commands and parameters set by digital platforms. One implication of this shift to examine is how it impacts traditional expectations of reading. Historically, communities of readers have approached texts with the anticipation of interpreting representations that reflect intentional human thought. In contrast, algorithmically generated inferences challenge these expectations by presenting outputs based not on human creativity or intent but on calculated data responses. Given that the traditional metrics for assessing the truthfulness or falsity of a text—criteria deeply ingrained in reader response theory—are upended, readers are left to navigate a new set of expectations about what constitutes true or false writing in an era where the lines between them are increasingly blurred by algorithmic mediation. This change in the nature of text from human-driven representation to data-driven inference complicates the perceptual strategies readers have long relied upon, forcing a reevaluation of how texts are approached and understood.

Shifting technologies necessarily catalyze a transformation in the capacity, necessity, and locus of trust assignment. In a world where text is created by human minds, trust in language is linked to the consciousness and attentive effort of the writer. However, in an algorithmic landscape, where outputs are based on probabilities rather than direct representations, such traditional assignments of trust become untenable. This evolution has profound implications for various literary genres, including journalism, where readers inherently assume the presence of a mindful, intentional author behind the text. This assumption underpins the very structure of copyright law, which has historically governed the relationship between the identity of the author and their creations, predicated on the belief in a conscious entity capable of original thought and intent. After all, if there is no conscious mind behind the creation or action, to whom do we attribute responsibility and ownership? For instance, when a self-driving car is involved in an accident, the challenge becomes identifying who is liable, and this highlights the complexities that arise when traditional concepts of ownership and responsibility are applied to decisions made (or outputs generated) by non-conscious entities. The absence of any conscious mind behind these technologies challenges our conventional understanding of copyright and responsibility, prompting a reevaluation of what constitutes an author or liable party—and therefore a locus of trust—in the age of automation.

Since the 19th century, our trust in automated systems has largely been systemic, located at the level of broad societal functions. Trust, fundamentally, underpins governance, evinced by the role of currency in an economy (a stable currency reflects the successful management of trust by the government). We routinely trust ATMs to dispense the correct amount of money and believe that our bank account numbers reflect actual values, embodying our trust in the abstract yet concrete value of money. However, when it comes to automated writing, we encounter a

significant hurdle in the sense that we have not reached a point where we can trust the abstract values produced by such systems as “meaningful” in the same concrete way. Our skepticism arises because automated outputs lack the human elements of consciousness and intention, upon which our traditional theories of meaning are founded. This disconnect challenges us to reconsider our concepts of meaning and trust in the context of increasingly autonomous technologies. It is therefore expected to consider whether outputs devoid of human oversight can truly possess, or convey, meaning in the way we have historically understood it.

The concept of exchange of value has traditionally been viewed as a uniquely human endeavor, deeply intertwined with the production and recognition of meaning. While we recognize that animals have systems of communication, we generally do not attribute to them a reflective relationship with meaning (or, at least, the kind of relationship with meaning that characterizes human language and interaction). Unlike animals and plants, which do not engage in systematic exchanges of value, humans are now confronted with the novel scenario of exchanging value with automated systems. While we often ascribe meaning to interactions with animals despite their lack of intrinsic meaning-making, we now must determine how to interpret and value exchanges with entities that operate without consciousness or intent.

3.5 Epistemological and Phenomenological Assessments

The political ontology of algorithms introduces complex challenges to epistemological assessments of algorithmic media, particularly concerning the theory of knowledge—what we know and how we come to know it. In the domain of automated news production, in which the use of large language models is increasingly prevalent, understanding the nature of knowledge production is crucial. Key questions arise: How can we ascertain the truthfulness, neutrality, or

factual accuracy of news produced by AI? What are the sources from which AI derives its knowledge, and how do these sources influence the authenticity and accuracy of the information it generates? These considerations are fundamental aspects of the central questions to this study.

Scholars and critics who analyze the social nature and functions of algorithms generally agree that they are not neutral. The prevailing methods for addressing the impartiality of algorithms fall under what might be termed “epistemic reformist” approaches. These critiques, of which contemporary examples abound, reflect how algorithms, far from being neutral, embed and propagate existing human values, biases, and decisions. Safiya Noble (2018) emphasizes this in her work, pointing out how algorithms can perpetuate dynamics of power and inequality through automated technologies of social evaluation. Similarly, Colin Koopman (2019) discusses the “infopower” of algorithms in *How We Became Our Data*, demonstrating how they can enforce discriminatory practices like redlining through the subtle formatting of information. Frank Pasquale (2015)’s notion of “black boxes” in technology critiques the opacity of algorithms, which obscures their functioning and challenges their trustworthiness. Ben Green (2020)’s analysis of criminal justice risk assessments calls for a shift away from “colorblind proceduralism,” advocating for a structural approach to reform that addresses deeper systemic biases. Together, these perspectives underscore the epistemic reformist view of algorithms, which serves foremost to challenge the assumption of their neutrality and call for a reevaluation of how they are designed and deployed in society.

Rethinking epistemic frameworks is essential when addressing AI-generated content, particularly given the ontological differences between traditional representational media and the virtual, non-representational characteristics of algorithmic dispositifs. These differences present unique challenges for epistemic analysis, which has traditionally relied heavily on

representational structures to ascertain truth and meaning. Unlike representational media that aim to accurately mirror reality, the algorithmic dispositif operates through virtual processes such as modeling, capturing, tagging, sorting, and rendering. These processes are not concerned with direct representation or accurate reproduction; rather, they focus on adaptive thresholds and probabilistic calculations. Therefore, conventional epistemic analysis, which is predicated on the accuracy and fidelity of representation, is insufficient, and the reliance on adaptive and probabilistic methods by algorithms calls for a new approach to understanding how knowledge is generated and represented in the digital age. Ultimately, critical approaches to AI-generated content must move beyond representational assumptions and adapt to the complexities introduced by these advanced computational processes.

Phenomenological considerations provide a critical augmentation to ontological and epistemological approaches when analyzing the impact of algorithms. While epistemological approaches focus on the knowledge and beliefs that people hold about algorithms, investigating how these understandings shape perceptions of algorithmic consequences, phenomenological approaches delve into the lived experiences of individuals interacting with these systems. For instance, epistemological analysis might reveal that people who perceive algorithms as biased are more likely to view their consequences negatively, whereas those who believe algorithms to be objective might see their impacts more positively. In contrast, a phenomenological approach emphasizes the actual effects of algorithms on lived experience, independent of beliefs or knowledge about the technology. This perspective examines the tangible ways in which AI influences individual confidence, decision-making abilities, and sense of identity. By exploring these direct experiences, phenomenology offers a deeper, more human-centered understanding of

the implications of algorithmic interactions, highlighting the real-world consequences that these technologies have on everyday life.

Peter Polack's theory of paramediation offers a perspective that calls attention away from traditional ontological assessments of algorithms toward a phenomenological assessment of how individuals experience them. This theory advocates for moving past a sole focus on algorithmic bias towards a broader analysis of algorithmic domains, examining how these technologies affect everyday life and interactions. This approach is critical of the conventional strategy of algorithmic reformism in its tendency to propose new algorithmic solutions to existing design problems without addressing the fundamental political stakes involved in these design issues. The political ontology of the algorithm dispositif encourages a deeper recognition of these stakes, suggesting a need to look beyond mere technical fixes.

Polack's approach is more experientially focused than traditional approaches of epistemic reform, emphasizing activities, practices, and actions over ideational constructs, and is less concerned with the underlying structures of domination that often remain hidden beneath the surface of experience. This perspective does not depend on the cognitive expertise of the critic, but instead on understanding the activity as an embedded practice. When considering algorithms within their technical milieu and examining the forms of participation they enable or disable, evaluating algorithmic interactions within a theory of paramediation allows us to uncover how algorithms introduce new forms of relationality and articulate existing modes of association in specific lived contexts. This nuanced view challenges the notion that the social consequences of algorithms are fully captured by their logical structure, advocating for a broader understanding of their impact on human experience.

Applying epistemological and phenomenological analyses to algorithms in journalism provides a holistic framework for addressing the three research questions concerning the production and reception of automated news content and journalism. From an epistemological standpoint, understanding how journalists and readers perceive particular characteristics of algorithms illuminates the biases and expectations they bring to bear on their engagements with automated content. An epistemological analysis is crucial for discerning how user beliefs about algorithmic neutrality or bias influence their trust in and reception of AI-generated news. Simultaneously, a phenomenological assessment offers a deeper dive into the lived experiences of these stakeholders, revealing how the actual use of algorithms in journalism impacts their professional practices and daily interactions with news content. Together, these considerations provide a comprehensive view, recognizing both the conceptual beliefs and tangible effects of algorithms, thereby enriching the understanding of how automated news content is produced, perceived, and experienced in contemporary journalism.

3.6 ChatGPT Case Study

A case study of ChatGPT serves as an illustrative example of the ontological and phenomenological dimensions of AI-generated text, offering insights into how AI differs fundamentally from traditional conceptualizations of writing. Traditionally, writing is viewed as a transcriptional relation between thinking and writing, which constitutes a direct representation of ideas. In contrast, ChatGPT operates not as a direct transcription but as a probability engine, synthesizing output based on nuanced regularities derived from extensive datasets. This shift from direct representation to probabilistic output highlights a significant ontological shift in how text is generated. Unlike tools such as spellcheck, which also assist in writing but do not trigger

significant moral concerns, ChatGPT embodies more complex capabilities that extend well beyond simple text correction to generating coherent and contextually appropriate text based on learned patterns.

The fundamental structure of ChatGPT, like all large language models, is a neural network that processes numerical representations of data to generate coherent text. As explained in Chapter 2, neural networks are composed of node layers: input, hidden and output layers. These layers work together to process information and produce a response. Neural networks reflect the behavior of the human brain, allowing computer programs to recognize patterns and solve common problems. Because of this, neural networks are sometimes specifically referred to as artificial neural networks (ANNs) to distinguish it from the non-artificial neural network that exists in the human brain. Each node, or artificial neuron, serves as its own linear regression model. The weights of each node determine the effect the input node has on the output.

Large language models are trained on massive amounts of text data, and use the data on which they are trained to learn statistical patterns between words. Subsequently, this knowledge is utilized to predict the next word, one word at a time. In the context of ChatGPT, the application is undergirded by a model (GPT-4) that was trained on internet data consisting of hundreds of billions of words. The model is trained to predict the next token given a sequence of input tokens, and is able to generate text that is grammatically and semantically similar to the text that it is trained on. However, these models need proper guidance in the form of user prompts.

From a user perspective, interacting with ChatGPT requires minimal technical acumen; the interface is in many ways similar to a search engine in that it “reads” human input similar to how Google “reads” queries (Floridi and Chiriatti, 2020). User interaction with ChatGPT is done

via prompts--textual hints that guide ChatGPT about the context of the input and task at hand (see Figure 3.1 below). These user prompts correspond to the stage of fine-tuning in the machine learning process, in which the weights of a trained neural network are used as initializations for a new model being trained on data from the same domain (Howard and Ruder, 2018). When a user prompts the system, they are able to guide ChatGPT in knowing which of the neural pathways it should activate to perform best. Based on the massive corpus of Internet data upon which it acts, ChatGPT has neural pathways or “trees” in its knowledge base or brain -- the user prompts guide ChatGPT in choosing particular trees based on what it predicts to be the most likely text to follow from the text that the user provides. To facilitate this directly, a user can provide examples of successful completions of a prompt within the prompt itself. The more examples of successful completions (pairs of inputs and outputs) one provides, the better chance there will be to generate a proper completion for an unseen case (Mor, 2020). Given the ability of the user to fine-tune the model via prompting, the quality of the model’s output is more of a function of effective prompts rather than the execution of the natural language processing task itself.

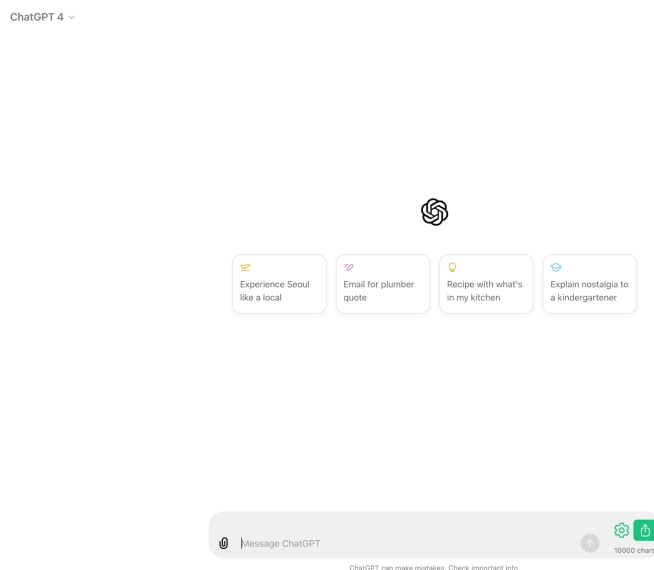


Figure 3.1: The ChatGPT (GPT-4) interface

A prompt in the context of the ChatGPT algorithm is essentially a database query: after it is passed in, ChatGPT gives a probability distribution of what the next word might be. The neural pathways or “trees” are sequences of probability distributions of what the next word might be given an existing input of a certain size of tokens or words. Two of the main search algorithms used in NLP to traverse these trees are greedy search and beam search (Daume et. al., 2006). Greedy search algorithms always take the locally best (highest probability) option at each stage. In the context of a language model such as GPT-4, this would entail selecting the most likely word to follow the previous textual input each time.

Given the similarities between prompting ChatGPT and querying a database, interaction with ChatGPT evokes similarities to interacting with Google: the initial user instinct is to ask the machine a question (Floridi, 2023). The corpus from which the ChatGPT system searches to produce its response is human-generated writing from the Internet. Search engines like Google have two main elements: a data center element that contains the information to be retrieved and a UX (user experience) element that emphasizes the aesthetic element of the user information retrieval process (Nudelman, 2011). Search engines with well-implemented UX design simplify the process of communicating intent to the data center in order to retrieve information from it. In terms of this dichotomy, the GPT-4 model itself is like the data center or data box, and the OpenAI API is the UX element. When faced with a user prompt input from the API that it must complete, GPT-4 combs over its corpus of data, consisting of words humans have written, in order to probabilistically determine its response.

Both large language models and the human brain procedurally turn ideas into words, but these processes manifest themselves differently. With regard to writing quality, GPT-4 performs well at some tasks and poorly at others. Its predecessor, GPT-3, performed less consistently than

human writing, often varying gender and personality descriptions erratically within passages. Issues such as semantic repetition at the document level, loss of coherence throughout passages, contradictions, and non-sequitur sentences and passages were the most notable weaknesses in terms of text synthesis for GPT-3 (Brown et al., 2020). Furthermore, GPT-3 output reflected difficulties with comparison tasks (e.g. if two words are used the same way in a sentence, or if one sentence implies another). It also performed poorly at grammar and common sense reasoning (Elkins and Chun, 2020).

GPT-3.5 and GPT-4 (the models underpinning Chat-GPT) outperform GPT-3 in the aforementioned tasks (Koubaa, 2023). GPT-4 in particular performs astonishingly well on a variety of standardized tests; for instance, GPT-4 significantly outperformed both human test-takers and prior models on the multistate bar exam (MBE), demonstrating a 26% increase over GPT-3.5 and beating humans in five of seven subject areas (Katz et al., 2024). However, GPT-4—and ChatGPT in general—still exhibits many of the limitations of its predecessors, including hallucinations: mistakes in the generated text that are semantically or syntactically plausible but are in fact incorrect or nonsensical. In other words, ChatGPT is prone to making stuff up and BS-ing. For example, it is sometimes unable to properly count the number of letters in a word (see Figure 3.2).

Where is the r in blueberry?



The letter "r" is not present in the word "blueberry." The word "blueberry" is spelled with the letters B, L, U, E, B, E, R, R, Y.

Figure 3.2: A ChatGPT Hallucination Example

The reason this happens is because the tokenization process of the semantics destroys the meaning of each individual letter by sometimes combining them. As mentioned in Chapter 2, the average token size is approximately four characters. So while ChatGPT can count the tokens upon which it has been trained, it cannot yet identify the specific letters within them that have been tokenized. Another, similar example arises when ChatGPT is prompted to write sentences with a certain number of words (see Figure 3.3 below).

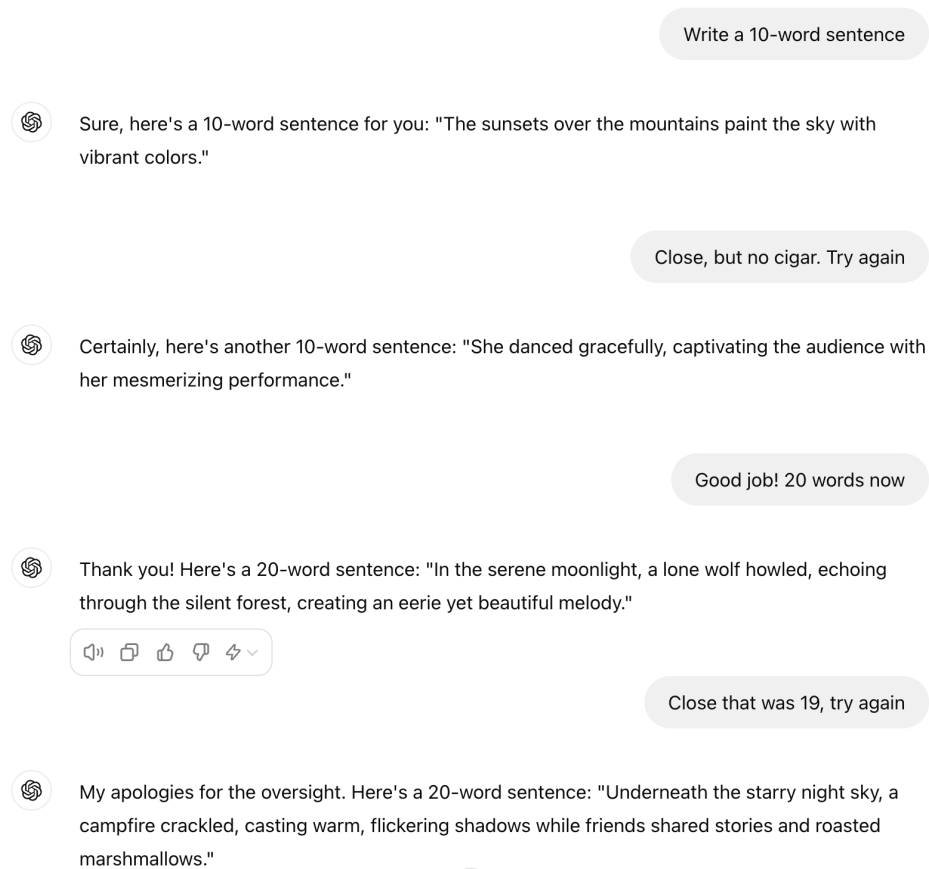


Figure 3.3: Another ChatGPT Hallucination Example

The problem of “hallucinations” belies a problem beyond just the AI model seeing something that isn’t there, reflecting the difference between mimetic representation and the algorithmic

ontology. Whenever an LLM generates an output, it is trying to produce human-like text, but not in the same way that humans are always “trying” to write well. It may well be true that the goal with writing is simply to make something that sounds good, but for the underlying procedures of generative AI, this goal is never directly tied to the world. When it goes wrong, it isn’t because it hasn’t succeeded in representing the world this time because it’s goal is not representing the world. It is not dealing in truths or falsities. For example, a hallucinating human whose vision was once useful in representing the world, and was good at this because of its connection to the world.

Beyond numeracy, the issue of unreliable outputs also includes factual problems with ChatGPT. On its own, it still requires frequent, massive retraining to be current with the news. GPT 3.5 knows nothing of 2022 and GPT-4 seems to know little of 2023 (Rosol et al., 2023). The reason for this is in part due to the cutoff dates in model training: a knowledge cutoff date is the specific point in time after which an AI model no longer has updated information. So when prompted with the task of writing a news article about a current event, it is to be expected that ChatGPT would produce an output with fabricated information. However, it is still trained on real-world data and can bring in relevant information, albeit outdated, into its text generation process. Figure 3.4 demonstrates an interaction with ChatGPT that exemplifies this concept. After the user inputs a prompt instructing ChatGPT to write a news article with a specific title, an output is generated that resembles the body of a news article. The inputted title was taken from an actual article from CBS News about anticipation of an atmospheric river that impacted

Northern California in early 2023 (Martinez, 2022).

Write a news article titled "Potential Category 3 atmospheric river to hit Northern California"



Northern California is bracing for a potentially powerful atmospheric river that could bring heavy rain and strong winds to the region. The National Weather Service has issued a warning for a Category 3 atmospheric river, which is expected to make landfall on Wednesday.

An atmospheric river is a narrow band of moisture in the atmosphere that can bring large amounts of precipitation to a specific area. This particular atmospheric river is expected to bring up to five inches of rain to the region and could cause flash flooding and landslides.

The strong winds that are expected to accompany the atmospheric river could also cause power outages and damage to structures. Officials are urging residents to secure loose objects and prepare for potential power outages.

"We are urging residents to take this storm seriously and to prepare for the potential impacts," said Mark Strudley, a meteorologist with the National Weather Service. "Make sure you have an emergency kit ready and be prepared to evacuate if necessary."

The atmospheric river is expected to last through Thursday and could bring much-needed relief to the region, which has been dealing with a severe drought. However, the potential for flooding and other hazards means that residents need to be cautious and take the necessary precautions.

Figure 3.4: A News Article Prompt

The output is structured in five short paragraphs of two sentences each: first, an introductory summary similar to that of a real news article; second, a definition of an atmospheric river and the danger of the particular atmospheric river in the article; third, additional general information about the anticipation of the atmospheric river. The fourth paragraph is the most notable here as it includes a quote from Mark Strudley, who ChatGPT labels as a "meteorologist with the National Weather Service." Interestingly enough, Mark Strudley is indeed the name of someone in real life who works in flood risk reduction in Northern California (see Figure 3.5). However, the quote in the article was never actually said by Strudley or anyone else, nor was Strudley ever interviewed about this particular atmospheric

river. However, he had been quoted in past news articles about meteorology (York, 2022).

ChatGPT was therefore able to cite Strudley in its article based on information from the corpus of past information, including similar news articles, on which it was trained, but fabricated a quote for the sake of following the instructions given to it in the user prompt.

Mark Strudley, PhD. · 3rd
Executive Director at Pajaro Regional Flood Management Agency
Santa Cruz, California, United States · [Contact info](#)

Figure 3.5: LinkedIn Header for Dr. Strudley

Ultimately, without proper guidance, ChatGPT can generate outputs that are untruthful, toxic, or reflect harmful sentiments. Even with that severe downside, ChatGPT itself is certainly useful—but only in a very structured way.

Chapter 4: Methods

Chapter 3 laid out the ontological, and epistemological, and phenomenological considerations necessary to assess both the production (writing) and consumption (reading) of AI-generated content. Given these considerations, this chapter overviews the specific methodological choices that underpin the study, first articulating the purpose of the study and the particular gaps that it fills in the literature on AI in news. This study employed a mixed methods approach of semi-structured interviews and surveys in order to explore diverse perspectives and uncover relationships that exist between the intricacies of unanswered questions regarding the production and consumption of automated news content. In this chapter, I will articulate my rationale for using a mixed methods approach and provide a detailed account of my study design, data collection, and analysis procedures.

4.1 Purpose of the Study

The interviews are designed in part to explore how journalism itself is responding to the novel ontological phenomenon of automated text generation. The problems that algorithms pose as political entities with powers of governance are ones with which to confront the world of journalism. Elements of this confrontation include the priorities of journalistic practitioners: what they want, what they think about this new medium, and accounts of their experiences with the artificial intelligence tools that facilitate automated text generation. In addition, it is crucial to consider the extent to which the adoption of new algorithmic newsroom technologies would potentially alter the way in which they write articles, bylines, and other content. The autonomy of practitioners within the realm of journalism also warrants examination in terms of how

newsrooms balance journalistic and algorithmic labor in the areas of reporting, editing, and curating.

Furthermore, the interviews are designed to understand diverse perspectives of a variety of individuals involved in the production of automated news content. These perspectives can offer insights into the varied experiences of those who operate at the intersection of technology and journalism. This category also includes new entrants to the field, or “strangers to the game” of journalism, who possess the potential to disrupt news organizations (Belair-Gagnon and Holton, 2018). New roles emerging at news organizations may be “disrupting news production through advancements in digital and social media” and might focus on “emerging technology meant to supplement or complement journalists’ work” (Belair-Gagnon and Holton, 2018: 75). Such perspectives are of critical importance to the field of journalism as well as the public due to the contemporary relevance of text generation models like ChatGPT. Interviewing journalists in addition to practitioners with non-traditional roles in journalism can uncover how the contributions of these practitioners differ from each other, and how they perceive their involvement in the journalistic process.

Such perspectives are of critical importance to the field of journalism as well as the public due to the contemporary relevance of this phenomenon. The constantly evolving field of artificial intelligence merits continued research within the industries in which it is being integrated in order for society to better understand its effects. Contributing to the continually evolving corpus of academic research on the evolving roles of practitioners working with artificial intelligence both directly and indirectly is thus essential.

The surveys are designed to assess the consumption aspect of journalism: the perceptions of everyday audiences who read and engage with news content. In considering the evolving

landscape of news media, it's crucial to distinguish between news content and other content types for both media industry and administrative purposes. However, an intriguing question arises when we consider the psychological processing of this content by audiences. Does the categorical distinction between news and other forms of content significantly impact how readers perceive and process the information presented to them? This question invites a deeper exploration of the psychological underpinnings that govern media consumption. Similarly the distinction between AI-generated and human-authored news is a critical consideration from a media production standpoint. Does the explicit categorical distinction between AI and human authorship impact how readers perceive and process information? Yet, from a consumer perspective, the importance of this distinction may not be as pronounced. Readers' engagement with news content, their trust in its credibility, and the value they ascribe to it could be influenced more by the content's inherent qualities — such as accuracy, relevance, and comprehensiveness — rather than by knowledge of its authorship.

Departing from the canon of empirical studies that have attempted to isolate the impacts of both the content itself and the attribution of authorship, the surveys seek to dive deeper into the constructs underlying the metrics used in these studies themselves. Are metrics like credibility still relevant, and also, should they still be relevant? The way people describe news content that is explicitly written by AI will potentially illuminate new constructs and constellations of concepts brought on by novel content and source cues. Furthermore, the surveys compare perceptions of AI- and human-generated news content based on these new constructs, exemplifying the potential for these findings to augment existing empirical studies on news perceptions.

4.2 Research Questions

This study sought to answer the following research questions:

RQ1: How do individuals involved in production of automated news content describe their experiences?

RQ2: What are the evaluative criteria used by readers in their perception of automated news content?

RQ3: How do news readers in the United States perceive automated news content and human-written news content relative to the descriptors derived from RQ2?

In order to answer these questions, this study approaches the topic of automated journalism from a perspective rooted in the post-positivist social science paradigm through employing interview and survey methods. This study aims to answer RQ1 by using semi-structured interviews, RQ2 using a mix of open-ended and Likert-scale surveys, and RQ3 using Likert-scale surveys. My research questions are threefold and relate to both methodological components of the study: the interview component and the survey component. The interview component is designed to understand the perspectives of two different groups: news workers and experts. From these interviews, this study aims to gain insight into the perspectives brought to bear on AI in journalism. Through patterns and relationships uncovered in the interview component, the interview component of the study aims to illuminate, among other dynamics, the nature of AI-generated text; this study will then aim to complement these findings in the survey component, which is designed to 1) assess the evaluative criteria used by readers to perceive AI-generated news articles and 2) compare automated news production with

traditional human-generated news on these descriptors and criteria. Regarding the former goal, this component of the study aims to assess the extent to which, given the generative nature of automated news, new normative frameworks are needed for people to evaluate automated text production. Regarding the latter goal, this study aims to identify criteria within said frameworks that respondents attend to when reading news content.

The choice of both survey and interview methodologies for this study is informed by the paradigm of post-positivist research. This study employs a survey that involves empirical measurement of quantifiable survey responses to gauge respondent perceptions and beliefs, and the qualitative interview component consists of semi-structured interview questions. These elements and their distinctions reflect higher-level distinctions between positivism and interpretivism, the two dominant frameworks within which social science research is conducted. Whereas the purpose of positivism is to predict, control, or explain via the framing of general laws; the purpose of interpretivism is to understand and reconstruct through transferable findings (Pickard, 2013).

My choice of both quantitative and qualitative methodologies for the study I am conducting is made based on their suitability to answer my research questions. The methodological choices made in this study are a function of the aims of my research questions, which are to obtain descriptions of experiences with both the production and consumption of automated news content to ultimately assess the suitability of potential normative frameworks based on these descriptions. With regard to philosophical approaches that influence the research techniques employed in this study, these aims are influenced by both positivism and interpretivism; this study does not aim to achieve paradigmatic purity. The methodologies proposed for this study have been chosen to reflect the contrasting assumptions of these

philosophical traditions. As do positivism and interpretivism, both quantitative and qualitative methodologies have advantages and drawbacks. At their most basic level, the difference between the two methodologies is reflected by their conceptions of reality: quantitative methodologies are based on the assumption that social facts are grounded in an objective reality, whereas qualitative methodologies are based on the assumption of social constructions of reality (Gorman et al., 2005). This study attempts to examine both philosophies to highlight the most salient aspects of each with regard to the topic of automated news consumption and production.

The descriptors used to identify the focal points of my research inquiry are inherently tied to what becomes salient or relevant over the course of the study. Thus, in this emergent process, boundaries will be created and defined through my own actions and attentions as a researcher. For both interview and survey portions of the study, the boundaries were defined by the questions I deemed relevant to the overall goals of the process. My research questions are thus influenced by the exploratory nature of the study and my aim to create boundaries by following the story that emerges from the data. The issue of non-physical space is also relevant to the development and description of my research questions as it pertains to the complexities of interacting on a digital platform.

The mixed methods design is essential in fulfilling the overall purpose of the study itself: investigating dynamics of algorithmic news production and consumption in order to establish a theoretical framework within which to measure potentially novel criteria of automated news perception. In light of the nuanced and sometimes contradictory perceptions that readers have towards automated news, the survey component aimed to collect data on the evaluative criteria employed by readers when interacting with AI-generated content, and use these criteria as metrics with which to compare AI- and human-generated content. Understanding these criteria is

vital for news organizations as they navigate the complex dynamics of integrating AI into their journalistic practices, balancing technological efficiencies with reader trust and acceptance. Thus, RQ2 emerges as a necessary question of the study, as the evaluative criteria readers use to perceive automated news content is of utmost importance. This question aims to uncover the specific factors that influence reader perceptions, offering insights into how automated news is received. New potential factors are therefore preferable to legacy metrics for empirical studies comparing AI-written to human-written journalism in terms of their emphasis on the potential unique aspects of AI-written content to which readers attend. RQ3 therefore necessitates the collection of data to test these new potential factors. The answers to these questions are crucial for shaping the future deployment of generative AI tools in journalism, ensuring that the content produced not only leverages the advancements in AI technology but also resonates positively with intended audiences.

4.3 Study Design

4.3.1 Interview Data Collection Overview

For the interview phase, I interviewed participants from two groups, news workers and experts, as a means of evaluating experiences with automated news production. Due in part to the constraints brought on by the COVID-19 pandemic, all interviews were conducted over Zoom. Despite this limitation, conducting interviews virtually offered several distinct advantages: convenience and flexibility in terms of scheduling and joining meetings being foremost among them. Additionally, the cost-effectiveness and wide geographic reach afforded by virtual interviews also augmented their appeal. Lastly, virtual interview settings can provide a

sense of anonymity, which can encourage more candid discussions and make participants feel more at ease in their own environments (Olliffe et al., 2021).

The interview phase ultimately sought to answer RQ1 through the collection of data on how people experience the process of producing news with artificial intelligence, and consuming news generated by AI, from an interpretivist perspective. The term process is used here to refer not just to the process of writing an article but also the process by which information is obtained as well as the design and development of the algorithms used to write the article. To that end, data for the interview component was collected using a sensemaking metaphor as a guiding concept. The sensemaking metaphor is part of a broader approach that “conceptualizes messages not as things to be gotten, but as constructions that are tied to the specific times, places, and perspectives of their creators” (Naumer et al., 2008). When individuals produce automated news content, gaps in knowledge emerge through the presentation of headlines, which are then satisfied through the construction of linkages, or bridges, between discrete information (Chen et al., 2015). The questions in the interviews thus pertained to the themes of sensemaking, trust, and knowledge. Interview questions can be found in the Appendix.

4.3.2 Survey Data Collection Overview

The survey component consisted of a three-phase methodology designed to elicit reader perceptions of AI-generated news. Phase 1 of the survey sought to gain a qualitative understanding of how people describe automated news content. After obtaining long-form impressions from respondents, individual terms were isolated and used as variables for Phase 2. In Phase 2, reader rankings of articles based on said variables formed the basis for the exploratory factor analysis, in which connections between variables were identified to extract

factors underlying reader perceptions. Phase 3 used an experiment to compare how news readers in the United States perceive AI (ChatGPT)-generated and human-written journalism in regard to the underlying factors obtained in phase 2.

The methodological choices for the three-phase survey methodology in this study were strategically designed to capture a comprehensive understanding of reader perceptions of AI-generated news. In Phase 1, the focus on qualitative, long-form survey responses was crucial for gathering in-depth, descriptive insights from readers. This approach allowed participants to freely express their impressions and thoughts, providing a breadth of perspectives on AI-generated content. The qualitative nature of this phase was instrumental in capturing the nuanced and varied ways in which readers perceive and articulate their experiences with automated news, beyond the constraints of predefined response options.

Phase 2, building on the groundwork laid in Phase 1, involved a more structured approach with a questionnaire eliciting Likert-scale rankings. This phase was designed to quantify the descriptors obtained from Phase 1, allowing for a systematic, measurable comparison of reader perceptions across different AI-generated articles. The choice to conduct an exploratory factor analysis on the correlations between these ranked variables was a deliberate one. It enabled the identification of underlying factors that influence reader perceptions, moving beyond individual descriptors to reveal broader patterns and themes in how AI-generated news is evaluated. Phase 3, based on the aforementioned exploratory factor analysis, also used a questionnaire eliciting Likert-scale ratings to assess perceptions of AI- and human-generated news on the factors obtained in Phase 2.

The AI-generated articles were generated by ChatGPT, using human-written prompts to instruct the model on what to write. The goal of the prompting process was to instruct ChatGPT

to generate the entire body text of a news article by giving it the title of a real article from Reuters written by a human. Prompts were crafted to be as short as possible while clearly communicating their intent to ChatGPT. As an example, the prompt given to ChatGPT to generate the article used in Phase 1 of the survey is shown below in Figure 4.1. Each of the prompts is shown along with the ChatGPT output in the Appendix.

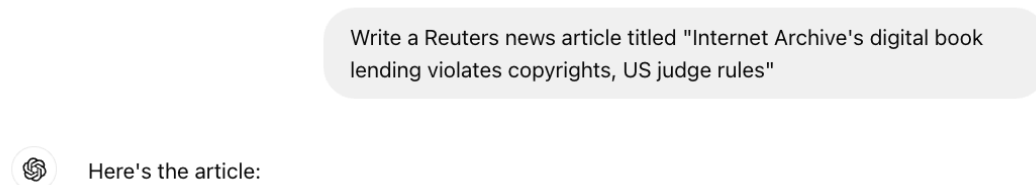


Figure 4.1: Prompt for Phase 1 survey

4.3.3 Interview Participant Recruitment and Data Collection Protocol

For the interview component, 13 participants were included— eight news workers and five experts. Participants included individuals from both the United States and the United Kingdom. As previously mentioned, these groups—news workers and experts—were chosen due to their proximity to journalistic practice, artificial intelligence practice, or some combination of the two. The group of news workers includes journalists, managers, startup founders, and engineers whose work relates to journalism practice: either working for a newsroom directly or licensing software or AI tools to newsroom clients. Experts spanned the domains of industry, technology, and policy.

Qualitative research is an iterative process, and as such, data collection and initial analysis for semi-structured interviews are often conducted concurrently. Preliminary analysis of interview data iteratively informs data collection and sample size. Contrastingly, in quantitative

research, a sample size is calculated in the planning phase of the study. In qualitative research, the sample size is considered to be sufficient when there are no new codes or themes generated from data collection. When such a threshold of sufficiency is reached, the research is commonly said to have achieved theoretical or thematic saturation (Adeoye-Olatunde and Olenik, 2021).

Researchers often draw on existing reports of saturation to estimate a sample size prior to data collection, which suggest between 12 and 20 interviews per category of participant. However, researchers have reported saturation with samples that are both smaller and larger than this (Knott et al., 2022). Correspondingly, this study reached thematic saturation after the completion of the aforementioned 13 interviews.

The journalist group was deemed relevant given their insights on how AI impacts their work and how it affects the ways in which they reach their audiences. Managers were deemed relevant due to their experience with shifting business models and innovation priorities due to technological innovation. Engineers were deemed relevant due to their use of AI in daily practice as well as their understanding of user dynamics. For the experts group, all three domains (industry, technology, and policy) were relevant to certain aspects of the interview questions and overall research methodology.

Snowball sampling was used to select interview participants in each group. Snowball sampling leveraged the networks within the respective communities within each group to identify potential respondents who would have potentially been difficult to reach through other sampling methods. Furthermore, respondents were generally more willing to participate if referred by a trusted colleague or acquaintance. Given the sensitive nature of the topic, the process of respondents being referred by initial interviewees was useful in terms of overcoming the barrier to entry of reluctance to participate due to confidentiality concerns. Lastly, snowball

sampling was utilized due to its allowance for adaptation of sampling strategies after learning more about the population. As interviews progress, researchers using snowball sampling can refine their focus based on emerging themes and recommendations from respondents (Sadler et al., 2010).

I invited prospective respondents to participate via email. In the introductory email, I provided my name, affiliation, contact information, and a brief introduction to the topic and background of my research. The introduction stated that I am studying artificial intelligence and journalism and included a short summary of the interview process, including a 30-45 minute virtual interview. I subsequently asked the respondent whether they would be interested in participating in the interview, which I noted would be audio and video recorded. If the respondent expressed interest in participating in the interview, I sent a follow-up email with an interview information sheet documenting the study procedure and scheduled a Zoom meeting with them via my UCLA Zoom account. Upon meeting with the participant, I re-introduced the background and context of the study, asked if they had any questions about the information sheet and study plan that I had sent to them via email, and obtained their oral consent to participate in the research study prior to starting the interview.

4.3.4 Survey Participant Recruitment and Data Collection Protocol

For the survey component, participants were recruited via online platforms for each of the three phases. Each survey was designed on an online platform as well. Table 4.1 below summarizes the design of each survey phase, the platform used to design, pilot, and/or conduct the survey, the stimulus or stimuli used in the survey, the sample size of the survey, and whether

or not the survey informed participants of the authorship (ChatGPT or human) of the stimulus or stimuli.

Phase	Platform(s)	Stimulus/Stimuli	n	Informed of authorship?
1	Amazon Mechanical Turk (pilot); YouGov (main)	One article written by ChatGPT (GPT4)	50 (pilot); 100 (main)	Yes
2	Qualtrics (design); Amazon Mechanical Turk (pilot/main)	Three articles written by ChatGPT (GPT4)	100 (pilot); 261 (main)	Yes
3	Qualtrics (design); Amazon Mechanical Turk (main)	Three ChatGPT (GPT4) and three human articles	153 (main)	No

Table 4.1: Survey Phase Design Summary

Phase 1

In Phase 1, participants were asked to detail their impressions and thoughts regarding an AI-written news article using descriptive adjectives. Specifically, the survey asked each participant to “List the thoughts that come to your mind after reading the article” and “List 2-10 adjectives describing the article” This phase was first designed and piloted in March 2023 on Amazon Mechanical Turk (MTurk), an online platform widely used for distributing surveys and tasks that require human intelligence. For the pilot survey, responses were gathered from 50 MTurk workers. Following a successful pilot, the final version of the survey was designed and distributed through YouGov, an international research data and analytics group, in April 2023. A larger sample size for this stage was sought, resulting in responses from 100 participants.

Amazon Mechanical Turk (MTurk) and YouGov, the two platforms utilized in this study, each offer distinct merits for online survey research, contributing to the robustness and reliability of the findings. MTurk, known for its diverse pool of users, is particularly advantageous for

piloting surveys due to its accessibility and the quick turnaround of responses. This platform is ideal for initial stages of research like Phase 1, where the primary goal is to gather a broad range of qualitative responses in a relatively short time frame. MTurk's user base, characterized by a wide range of demographics, provides a heterogeneous sample, essential for capturing a variety of perspectives and ensuring that the survey is accessible to a diverse audience.

On the other hand, YouGov specializes in providing access to a more targeted and often larger sample of respondents. This characteristic was particularly beneficial for the final version of the survey, as it enabled the recruitment of a larger and more representative sample. YouGov's strength lies in its ability to reach specific demographic groups, enhancing the generalizability of the findings. The platform's robust analytics and data cleaning capabilities further ensured the reliability and validity of the data collected in Phase 2. By leveraging the strengths of both MTurk and YouGov, this study benefited from the agility and diversity of MTurk's respondent base in the pilot phase and the depth and representativeness of YouGov's panel for the final survey, ensuring a comprehensive and nuanced understanding of reader perceptions towards AI-generated news content.

Phase 2

The survey for Phase 2 was designed on Qualtrics and deployed on Amazon Mechanical Turk. Qualtrics, as a survey design platform, offers several advantages that are particularly beneficial for social science research and were instrumental in the planning of this experiment. Firstly, Qualtrics provides a highly intuitive and user-friendly interface for designing surveys, which is crucial in crafting detailed questionnaires like those used in this study. Its flexibility in question design allows for a range of response types, from open-ended qualitative questions to

structured Likert scales, facilitating the collection of diverse data types that are essential in social science research.

Another significant advantage of Qualtrics is its robust data collection and analysis tools. The platform ensures the accuracy and integrity of the data collected, which is paramount in research settings. Real-time data tracking and advanced analytical tools allowed for immediate insights into response patterns, enabling quick adjustments if needed. Additionally, Qualtrics' compatibility with various data analysis software streamlined the process of exporting data for more detailed statistical analysis, such as the exploratory factor analysis conducted in this study. Furthermore, Qualtrics' strong emphasis on data security and confidentiality is critical in social science research, where sensitive information is often collected. The platform's adherence to stringent data protection standards ensured that respondent information was securely managed and processed, maintaining the ethical integrity of the research. Qualtrics' combination of ease of use, advanced functionality, robust data analysis capabilities, and strong security protocols made it an ideal choice for conducting the complex and multi-faceted survey required in this experiment.

A pilot survey for Phase 2 was conducted in June 2023 with a sample of 100 respondents. The pilot study allowed for the testing of the design and format of their survey, including the wording and order of questions, the layout of the survey, and the functionality of interactive elements. Following the pilot, an anonymous survey was carried out between August and October 2023 as part of the main Phase 2 study. For this survey, a sample of 261 respondents was obtained. This number was deemed adequate given the 26 variables being measured in the data set, and the general rule of thumb for adequate EFA sample sizes is 10 subjects per variable (Nunnally, 1978). Upon commencing the survey, each participant was shown three articles. To

write each of the articles, ChatGPT (GPT-4), a large language model developed by OpenAI, was prompted with the title and source of an existing, human-written article in Reuters with the same title. The output was not edited in any way. Participants were informed that each of the articles were generated using ChatGPT.

Each of the articles covered a different topic to provide a breadth of content for measurement: politics, finance, and technology. After each article, participants were prompted with these instructions: In the following question, for each word, please rate how well the word describes the article above, from “describes very poorly” (1) to “describes very well” (5). Participants were presented with the 26 descriptors to rate at the end of each article. At the end of the survey, each participant was given a unique survey completion code to submit via MTurk. After data collection was finalized, exploratory factor analyses (EFA) were conducted to scrutinize the differences in factor structures, thereby shedding light on potential variations in perceptions of stimuli. The full Phase 2 (descriptor ranking) survey is located in the Appendix.

Phase 3

Similar to the survey for Phase 2, the Phase 3 survey was designed on Qualtrics and deployed on Amazon Mechanical Turk. The survey was carried out between November and December 2023. For this survey, a sample of 153 participants was obtained. Participants were asked to read three news articles, and were instructed as follows: “Please read the articles and answer the subsequent questions.” Unlike in the previous two phases, participants were not informed of the authorship of the article. The purpose of this was to directly compare the articles themselves rather than including any authorship cues. A 2 (authorship - between subjects) x 3 (topic - within subjects) factorial design was used. Each participant assigned to the Human group

read 3 articles written by Reuters, each on a different topic: finance, politics, and technology. Participants were asked after each article to “rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5). The ranking sliders included clarification of terms to ensure respondent understanding. Survey questions are presented in more detail in the Appendix. Below, in Table 4.2, are the word counts for each of the articles. The average word length was similar—for the AI-generated articles it was 386.67 and the average word length for the human-written articles was 379.67; however, more variation was observed in length among the human-written articles.

Article Authorship	Article Topic	Article Length (# words)
AI	Finance	391
AI	Politics	407
AI	Tech	362
Human	Finance	536
Human	Politics	364
Human	Tech	239

Table 4.2: Phase 3 Article Length

4.4 Interview Analysis Overview

Each of the 13 interviews lasted between 30-45 minutes. When the participant joined the virtual meeting room, I introduced myself, reviewed the interview information sheet, and asked them if they had any questions or concerns. Following an informal initial conversation, I asked each participant if they consented to audio and video recording, if they would like their data to be anonymized, and subsequently if they were comfortable beginning the interview at that point.

Three of the 13 participants' data was anonymized. As such, in the Interview Analysis chapter, they are referred to by their profession and their regional geographic location e.g. "a U.S. east coast data journalist". In the case where participants requested their data to be anonymized, I recorded audio only. After each interview, I took summarizing notes and transcribed the audio recording. I used Zoom's recording function to record the interviews and saved the resulting files directly to an encrypted password-protected hard drive.

To design interview questions and analyze participant responses, I followed Joseph A. Maxwell (2012)'s "interactive approach" to qualitative research design. This first entailed aligning the questions with my research goals – RQ1 in particular. Next, this involved creating open-ended questions that are clear and easy to understand, and free of jargon or overly complex language. I also accordingly applied Maxwell's methodology to the sequencing of questions, starting off with general questions to make the participant comfortable before moving to more specific questions, making it easier to obtain detailed responses about the experiences of the participant with AI and journalism in their particular role. I also made sure to word questions neutrally, and in a non-leading manner, while still remaining flexible and prepared with probing and follow-up questions to explore participants' answers in more depth (Pickard, 2013).

Over the course of the interviews, I identified common themes and theoretical issues for further investigation, and I took notes on participant responses. I began by coding participant responses to group them into thematic responses. Labels were assigned to chunks of texts from interview transcripts based on their content and meaning. During the coding process, memoing was utilized to capture my thoughts about the codes and categories, helping to explain why certain codes were chosen and how they relate to each other. By taking notes about the data being collected, I reflectively and dynamically tracked my analytical thinking and decisions

across the interview coding process (McGrath, 2021). Participant responses were analyzed in terms of their correspondence to the fundamental assertions identified in Chapter 3, including the notion that AI-generated content is a fundamentally novel entity as well as more broad notions of AI and cognition. This analysis therefore enabled me not only to use the organized data to answer my research question RQ1, but also to examine the relationship between participant responses and the theoretical considerations underpinning the study.

4.5 Survey Analysis Overview

The responses from Phase 1 of the survey, after being cleaned, refined, and narrowed into a suitable list of one-word descriptors, informed the variables rated by participants in Phase 2 of the survey. This process is discussed in more detail in Chapter 6. In the exploratory factor analysis of Phase 2, two distinct factor rotation methods were employed: varimax and oblimin. Varimax is an orthogonal rotation technique that simplifies the interpretation of factors by maximizing the variance of squared loadings of a factor on all the variables in a factor matrix. This method assumes that the factors are uncorrelated (orthogonal to each other) and thus focuses on clarifying which variables load highly on each factor. Oblimin is an oblique rotation method that allows for correlations between factors. This is more realistic in many social science contexts where underlying constructs are often interrelated. The oblimin method can provide a more accurate representation of the data when factors are expected to be correlated.

The decision to include results from both varimax and oblimin rotations in the study was driven by the objective to gain a comprehensive understanding of the data. While orthogonal rotations like varimax simplify the structure and interpretation of factors, they might not always represent the true underlying relationships in the data, especially when factors are correlated.

Oblique rotations like oblimin, although potentially more complex to interpret, can offer a more realistic picture of these inter-factor relationships. By examining the results of both rotation methods, the study leveraged the clarity and simplicity of orthogonal rotations and the realistic representation of factor correlations provided by oblique rotations. This dual approach allowed for a more nuanced and thorough exploration of the factor structure of reader perceptions of AI-generated news, ensuring that the findings were both interpretable and closely aligned with the structure of the data itself.

Hypotheses were formed after the completion of Phase 2 based on the factors obtained as to the expected effect of authorship on said factors. For the data obtained in Phase 3, a two-way, repeated measures ANOVA (analysis of variance) was performed on each dependent variable. Authorship (AI vs. Human) was tested as a between-subjects factor given that different subjects were exposed to different levels of this factor. Topic (Finance vs. Politics vs. Tech) was tested as a within-subjects factor given that each subject was exposed to all levels (saw all three articles) of the factor. This study incorporates authorship as a between-subjects factor and topic as a repeated within-subjects factor, which lends itself to a two-way repeated measures ANOVA. The design examines the main effect of authorship (how outcome variables differ between AI and Human articles regardless of topic), the main effect of topic (how outcome variables differ across the three topics, regardless of authorship), and the interaction between authorship and topic (whether the difference in outcome variable ratings between AI and Human articles varies by topic).

Chapter 5: Interview Analysis

In Chapter 3, we saw how ontological, phenomenological, and epistemological considerations were all important in foregrounding the research questions of the study and helping to understand AI-generated news production as a phenomenon. The dispositional powers of the algorithm and its virtual mode of existence position algorithmically-generated media in a category distinct from mimetic media, and as such, an ontology of AI-generated textual outputs merits an analysis beyond the representational nature of human-written language. But how do those actually involved in news production reckon with generative AI as a technology? And if practitioners are gaining experience with the technology in practice, how influential is it in their day-to-day operations relative to the other tools they use? The extent to which it affects their experiences—on both a systemic and individual level—is essential to obtaining a broader command of the impact of AI—and large language models in particular—on the journalism industry, and thereby on the public sphere.

With regard to the 13 semi-structured interviews conducted to address *RQ1: How do individuals involved in production of automated news content describe their experiences?*, the following analysis is organized into coded themes that were identified during analysis of interview data:

1. AI as a Tool
2. Innovation
3. Business Models
4. Audience
5. Trust and Transparency
6. Guardrails

7. The Human Element
8. Value Alignment

After reviewing and categorizing the interview data based on thematic elements, I arranged these themes to pinpoint those that were significant or common across all respondents. Given that the interviews were with both news workers and experts, in the analysis below, I specify which of the two categories are involved in discussing the particular theme in question.

5.1 AI as a Tool

Both news workers and experts emphasized the idea of using AI as a tool in journalistic practice. Respondents elucidated how AI is impacting newsrooms at an individual level, and the nuanced multitude of ways in which AI (and automation more broadly) serves as a journalistic tool across the news production value chain. In terms of the impact of artificial intelligence on news, journalist and London School of Economics professor Charlie Beckett identifies the twofold—systemic and individual—level at which it takes place:

So we've got our business model level change and the structural change... in terms of use cases, I think it's everything. And it's going to be at a systemic level, in the same way that people used AI to scrape data to organize paywalls, or generative AI will be able to do that even better. But then there's going to be the individual use of it, AI as a tool. So transcription, translation, wiki, formatting, all that kind of stuff.

Beckett highlights an important distinction here between the impact of AI at the systemic or organizational level in terms of reshaping business models, and at the individual level in terms of

the professional tasks journalists perform to execute their day-to-day responsibilities—the realm in which the theme of AI as a tool becomes salient.

Beyond the distinction between systemic and individual impacts, it is important to consider the distinction between generative artificial intelligence, a subset of the field, and artificial intelligence more broadly as it pertains to assisting reporters. A data journalist at an eastern U.S. news outlet elaborates this distinction more specifically:

The use of AI for sort of data journalism stuff and NLG for articles are two completely separate phenomena, united by the phrase AI. Yeah, nothing to do with each other. Besides they use computers and some of the people who enjoy fiddling with one enjoy fiddling with the other, but they're pretty distinct phenomena.

The respondent brings up two key ideas here: first, invoking the concept of natural language generation, or NLG, as a separate method from other machine learning processes, and the role of AI more broadly to inform and assist journalists throughout the reporting process: not just producing language, but serving as a tool for practitioners to pursue a multitude of other ends. He continues:

...a lot of times the questions that we need to answer that are answerable, require us to gather our own data, because it's not about from the Census Bureau and categorize it or sort it or count it in ways that are complex. So to collect the data requires, often—often requires some substantial software engineering, and that sort of categorization. Sometimes you can make do with simple

1980s or 1990s style filtering and searching techniques. And sometimes you need machine learning and AI.

The evolution of data collection and analysis techniques for journalists is apparent to practitioners today, as evidenced by the respondent's illustration of the requirements that modern data journalists face. Furthermore, the respondent elaborates on the impact of using AI as a tool to answer journalistic questions:

...the goal is just answering the question that we need to know the answer to: how much sort of violent incitement and election fraud claims were on Facebook immediately before January 6, for instance? That's a question that I think is just one that we want to know the answer to. And so to gather the data that we needed in order to answer that question, we needed some AI.

Artificial intelligence as a tool for journalists to complete tasks, he explains, is—and has for a long time been—a crucial element of newsroom practice. Practical—and normative, according to the data journalist—uses of artificial intelligence in journalism range from verification, to information gathering, to assisting in the pre-writing phases, to other miscellaneous tasks. The data journalist continues:

...you want to use AI for things that are a lot easier to verify: this is a picture. This is a satellite image of an airstrip, for instance, this tweet calls for violence. That's sort of what mechanisms do another. Another instance, where I often use AI is to help me find examples where I'm looking for tweets that call for violence. It's not something you can really do with keyword searches. But

you can make an AI model to do that. And then you write about the examples. And you don't even mention the AI.

Key here is the idea that “you don’t even mention the AI,” because of the innateness of the tools involved in the journalistic process. Outside of the capacity to generate text, respondents have found the ability of AI to select, filter, process, and otherwise analyze data to be an essential and valuable task to the production of journalism. Journalist, professor, and legal expert Jared Schroeder highlights a recent example of this:

The New York Times, I think it was in December, did a report about how many 2000 pound bombs Israel had dropped in Gaza. And they used AI to look at satellite images, they trained AI to find the craters, to measure the craters. And that's how they figured it out. AI helped make the story possible. That wasn't possible otherwise.

Another crucial categorization of AI impact is between front- and back-end. Damon Kiesow, media expert and professor at the University of Missouri, explains that AI tools and their role in the journalistic production process are not always visible and understandable to the end user (reader):

I put these in different categories, there's AI as a tool, which is doing work at a scale that humans couldn't do easily, but also is back end work. And then there's the front end, consumer facing work, which is being done very often without cause or concern for any of the journalistic ethics.

Most news workers who were interviewed had some experience with artificial intelligence on the back end as a tool. However, experiences with generative AI (and its subcategory of NLG focused on text generation) work on the front end varied across respondents. Multiple participants, including some interviewed prior to the release of ChatGPT and similar user-facing NLG tools, possessed experience with automated text generation, but in a more structured capacity rather than prompt-based text wholly generated through deep learning. Journalist, professor, and legal expert Jared Schroeder articulates this distinction:

AP has been doing work with automating reporting for at least a decade. [NLG, natural language generation]-- those are not quite as powerful as what we're talking about, they were very much: you give us the information, and we will make a thing. But if you don't give us the information, we can't make a thing. But it does show that these types of tools can be used to automate certain types of journalism that don't require a lot of human intervention, like the daily stock reports, which are just so boring. But people want to know, it's valuable to an audience. It's formulaic.

Schroeder posits the example of the Associated Press using template-based natural language generation: a more structured and “formulaic” form of NLG where the creativity is limited to the variation allowed by the templates. As an instance of this brand of journalistic production, Schroeder mentions stock reports. Sports reporting is another realm of journalism that lends itself to this aspect of NLG. Jay Allred, the CEO of Source Media Properties (including three local news sites, a full service digital marketing agency, and an AI startup), recalls the early stages of

his efforts to build an AI product with which local newsrooms in his home state of Ohio could automated sports reporting:

It was very much built in a very experimental phase, there was no business plan. We were not envisioning it becoming a company with multiple products, or any of those things. We were just, what if we could cover every high school football game in the state of Ohio, automatically, while our editors were sleeping? Could we do that?

He tested his tool on a small subset of high school football games in Ohio, and describes the results of that test:

And you know, it worked, it created assets that publish on our site and happen automatically. And so then we just went through the iterative process of continuing, trying to make it better, and then thought maybe after about a year of testing, maybe we have a product here. Maybe there's a company, maybe other newsrooms would want this.

The advent of ChatGPT in November 2022 catalyzed heightened attention toward front-end, user facing applications of language generation techniques. But these have not been the only instances of language models being used in news production. Charlie Beckett contrasts recent advancements made in front-end technology with the last four years of NLG development on both ends:

...this last year has shown that it's a very exciting, dynamic topic. Perhaps the last four years, it's been less exciting, because the technology, the traditional machine learning, etc., natural language processing was very much something that was relatively inaccessible as a technology. With generative AI, it's the opposite.

In the experience of Ryan Restivo, a product director at Newsday and developer of the YESEO app for newsrooms, the use of machine learning and language generation techniques behind the scenes is much less glamorous from a public point of view than the text production abilities of a tool such as ChatGPT or Google Bard:

I was using the unsexy AI, the natural language processing tools that we use, that we still kind of call it that. But AI is ubiquitous for all these other things now. So yes, we do use a lot of machine learning to kind of extract those keywords in a story.

Restivo describes the AI used in journalistic production pre-ChatGPT as “unsexy,” pointing to the idea that consumer-facing technology is perceived as more valuable and appealing regardless of its overall utility compared to similar tools that facilitate back-end development and production. Restivo explained how he was able to apply the practical benefits of GPT models toward his tool, YESEO, an app that allows its users—particularly journalists—to brainstorm headlines:

...three quarters of the way through an eight month fellowship, all of these GPT things started flying around like crazy. So I had to rethink and reset. And then think about just, is there any

practical application here? I'm lucky, I was talking about a developer friend, and he told me how he was using GPT to teach him Ruby, and [...] TypeScript, even though those weren't the languages that he knew very well. And so that made me rethink: what could I do with the data that I already have that I'm testing with?

Restivo describes the experience of a “rethink and reset” that was common among many interview participants following the release of user-facing LLMs. This mindset enabled him to develop the YESEO app, a tool that uses language production but in a pre- or post-writing capacity that suggests headlines for journalists:

...that kind of brought me around to this idea of, well can we take what's in the story and kind of build ways to suggest things to people? And this wonderful Suggest Headlines button that people see so often, and it works out? People love the button, but they love it for many reasons. It's one of the reasons that really drives SEO, but it's also because they're looking for that shortcut...

Individual journalists are currently reckoning with the potential for generative AI models to write and publish articles, and the extent to which this can be done independently. However, when considering the role of AI as a tool, respondents emphasized that its value add is much clearer in assistive, back-end capacities. Joe Amditis, journalist and assistant director at the Center for Cooperative Media, explains that in his experience, understanding the true value of AI as a tool can be a knowledge gap at first for some newsrooms:

...one of the biggest gaps in understanding is that the real value there is going to be not in necessarily, well, yes, somewhat in prompting and understanding how this new way to communicate and relate to our computers, and tell it to do what we want, but in the ability to connect that, that processing and natural language capability to other aspects of their operations through, through basic, simple If This Then That style automation.

Amditis focuses on the “other aspects” of newsroom operations because he sees the direct value addition for newsrooms, but also because of his own experience with AI as a publisher and/or writer of content:

...another reason I focused on the back office stuff and internal users, whether it's onboarding, or it's style guide compliance, helping to write all texts doing things like this, is because I personally don't like it... I feel weird about the idea of it publishing stuff.

His hope for newsrooms and freelance journalists is that they recognize that AI can be safely integrated in a “back office” capacity, rather than relying too heavily on its text production capabilities to completely renovate their operations:

Whether it's from a freelancer that they're working with for the first time, or whether it's assisted by a robot. You don't just throw it out there when you get the copy, you still have to do an editorial process verification, transparency, fact checking, all that stuff. And I don't think and I don't think newsrooms are—I hope not all the good ones, at least—aren't just abandoning that

because the bot exists now. I think it's important for them to integrate that as part of their workflow, and in an assistive capacity.

Other respondents agreed. Charlie Beckett explains that AI has so far been ineffective when it comes to content creation:

There's a lot of friction, and there's a lot of disappointment. I spoke to somebody recently, who is retooling a whole kind of regional newspaper, in Germany. And some of the sophisticated content creation stuff just didn't work. It just needed so much rewriting that you might as well write it yourself... just wasn't good enough.

However, he emphasized that generative AI has proven to possess a high degree of utility outside of the realm of content creation, in terms of saving time and augmenting efficiency:

On the other hand, some brilliant summarization tools that were working really brilliantly, some really good subscription recruitment and tools were working really, really well, so it's going to be context specific. And the results are going to vary. Sometimes it's going to be nice to have transcriptions of interviews for journalists; it's nice to have and it probably will save some time in that you're a bit more efficient. Will it change your life? No, not really.

Beckett describes here how, while not life-changing, the impacts of generative AI across newsroom operations often work “brilliantly” when applied in the proper contexts. Jay Allred

acknowledges the limitations of fully generative output but is still interested in experimenting with it as it evolves:

We've been testing generative since ChatGPT came out in November of 2022. We immediately began to test it and as it stands right now, we think that we're getting close to being able to use generative in an output producing very guard-railed low risk, kind of an environment—we think we're getting pretty close to it, but we're not live with anything yet.

To illustrate the dynamic between the current value and limitations of generative AI tools, Amditis offers the following analogy of AI tools as metal detectors:

I like to treat these these tools as sort of metal detectors, as I mentioned in the article, where they're able to tell you if there's metal beneath the sand on the beach, but they can't tell you what kind of metal it is, why, how it got there, why it is valuable, who might find it valuable or interesting, who owned it before, the broader human context and and rich tapestry of narrative and human experiences is not contained in there, but it can be useful to point you in the direction of something that might be worth exploring more. And so that's how I tried it.

5.2 Innovation

For both news workers and experts, an essential concept of interest was innovation. Individuals in and out of newsrooms described the relationship between newsrooms and technology historically, and how that has informed both resistance and evangelism toward AI technology from news organizations. With regard to news workers, innovation is significant

because it drives economic growth for their organization and enhances competitiveness. Experts are interested in innovations as they pertain to the challenges posed by new technologies such as generative AI.

For both parties, the wave of interest in AI catalyzed in 2022 by the release of ChatGPT has led to both hype and significant change. Journalist and entrepreneur David Cohn elaborates on the nature of rapid advancements in AI and the extent to which these advancements are sea change versus merely hype:

...and I mean, it's a little bit of both [sea change vs hype] is what I would say, right? So I forget the name of it, but there's this general innovation curve... the expectations are that it's going to do everything in the world for you. And then inevitably, the technology doesn't meet that expectation, it drops. And then it actually comes back up and plateaus somewhere in the middle.

Cohn evokes the concept of an “innovation curve”, which explains how new ideas and technologies spread through cultures, to articulate how adoption of AI among newsrooms is not uniform. On the topic of generative AI in particular, Cohn describes how hype has shaped perceptions and expectations of disruption of existing technologies, and the importance of acknowledging that while it may not “do everything,” significant disruptive potential does exist:

I do think in its initial push, it is being overhyped, a little bit because it's not a panacea. It's not going to just instantly change everything. Or make journalists immediately make journalists... useless... you're still gonna need people. It's not going to do everything. That said, there is a non-zero chance that search is wildly disrupted, there is a non-zero chance that, certain types of

stories can be automated, you know, accurately and with confidence and definitely in the future, there's a non-zero chance that the user interface that most people use to get their news and information changes drastically.

Other journalists, including Joe Amditis of the Center for Cooperative Media, share some excitement about generative AI as a “different” technology. Amditis describes his “holy s***” moment that he felt when initially experimenting with ChatGPT:

*And I just finished reading Brandon Sanderson, his latest book at the time [The Lost Metal]. I'm a huge fan of his work and everything. And [the book] just came out like two weeks before, and I was like, let me see what this thing can do. So I started asking him questions about the book and the Stormlight Archive and the series and blah, blah, blah. And it was answering and talking about spoilers in this book that had just come out two weeks ago. Now at the time, I didn't know that it was largely making s*** up. And it was I was able to say, actually, no, this character doesn't actually do that. And he would say, oh, yeah, you're right. And it was just this moment where I was like, Holy s***, this is different.*

Intrigued by ChatGPT's capabilities, Amditis tested it by discussing Sanderson's novel, expecting accurate insights. Surprisingly, the chatbot responded with detailed information, including spoilers, but Amditis soon realized that some of the information was fabricated. Correcting the inaccuracies led to a realization about ChatGPT's conversational adaptability, sparking a moment of awe at the novelty and potential of the technology. Amditis's reaction highlights how the initial hype around ChatGPT's capabilities can lead to elevated expectations

about its accuracy and knowledge, showcasing the intriguing potential of AI to engage deeply and adaptively, even when its responses may initially contain inaccuracies.

Amditis also compares ChatGPT to previous technological advancement with which he has experimented in his life, namely Photoshop tools.

...I was doing Photoshop stuff in the early 2000s, making designs and sigs for Counter Strike. Every time Photoshop would release a new feature, I remember vividly the day they released content aware fill, or the healing Brush or the patch brush. And there was never any industry. I mean, I was 14. So I don't know about the industry. But there was never really among the communities that were working with designers and artists, there was never that kind of hesitancy where now, you need to label every time you've used a filter, or every time you've used a healing brush tool, or a spot correction or whatever. And we also were eager to use these new tools, and this time, it feels different.

His notion that “this time, it feels different” belies a sense of excitement and acceptance that traditionally greeted new technological features in software like Photoshop, contrasting with the caution surrounding the reception that generative AI tools such as ChatGPT are experiencing today.

And indeed, the technology is different, in the sense that generative AI is not the same as AI but rather a subset of the field. Insofar as generative AI is unique, it must also be considered in relation to past and future developments within the grander scope of artificial intelligence

research. Journalist, author, and professor Jeff Jarvis explains that experts who strategize around AI governance view generative AI in such a manner:

...the AI experts of the room don't really separate out generative AI from AI. To them, it's just a straight string, right. So in their view, the first stage of AI was analytical, call it, looking for anomalies in data and such, the second state is generative. And the third state is agentic. And so the question is, when will we trust the machine to be an agent on our own to go ahead and make that plane reservation for us, or write our taxes, or whatever. But I think that's where we also go next.

A computer scientist and generative AI startup founder in the U.S. shares the awe that Amditis has toward the technology. But similar to Amditis and other respondents, she emphasizes both the positive and the negative impacts that it forebodes toward society:

Yeah, so on generative AI, I think it's one of the most amazing inventions of humanity. And I think that there was gonna be a lot of positive, because it's super cool. And it's just so much positive possible here. But it's just really sad to see that there is unambiguously now going to be a lot of negative as well.

At the same time, practitioners in newsrooms are quick to remember that technology has always affected the news business in meaningful ways, and these changes often come in waves. Jay Allred makes note of this concept as he describes his own experience in newsrooms:

And it's important to note, historically, the technology has always changed. I mean, I have people in my newsroom that are old enough to have written stories on typewriters, so this is a big technological shift, for sure. But it's not the only time technology has affected the news business in a meaningful way.

Furthermore, Ryan Restivo underscores the fact that generative AI represents less of a leap than we might think, given the extent to which AI is already a relevant factor across many back-end newsroom processes that doesn't involve text production:

...my argument always has been: there's more under the hood, machine, like the broader AI term, NLP and machine learning and these kinds of things. There's a lot more that's already being done in newsrooms. So it's not like this is completely breaking new ground. This is just furthering that, so it's not like the leap of going to the internet. Like it was in the late 90s. It's like, already here, they just, there's just a little bit further. And maybe that makes them. Maybe it makes them a little less cautious.

Restivo brings up an important point regarding newsrooms in relation to this innovation: the dynamic between caution and eagerness to adopt. One of the risks of hype as a technology gains traction is premature adoption—if hype pushes technology into the market before it is fully developed, it can lead to consequences. Damon Kiesow, media expert and professor at the University of Missouri, emphasizes the idea that the response to the hype is what can cause problems more so than the hype itself:

The problem is less the hype around the technology, and more the the naive willingness of like news organizations to adopt any technology, which in their wildest dreams has any chance of improving efficiency, reducing costs, you know, speeding up processes, that kind of stuff, like questions about quality and trust, reliability, authenticity. Like none of those really come to mind, especially in the the many, at least half a dozen, dozen, sort of mini scandals we've seen in terms of people using generative AI bots to like, produce news content.

He goes on to explain that organizations that have followed the hype and prematurely adopted certain generative AI technologies did so not because of the technology itself, but because of factors within the organization that drove the particular decision-making process:

And it's got to do with this. This sort of insane drive for profit making and efficiency in the organization, because those that make these mistakes are not typically what you would look at as really, highly trusted sources of journalism. So it's easy to blame that on AI. But AI didn't force them to do that right? So it's not the technology.

Based on interview responses, on the spectrum between caution and recklessness, most organizations seem to be erring on the side of caution. Charlie Beckett sums up most newsroom approaches:

...so most mainstream, sensible news organizations are being relatively cautious. Some, perhaps overly, most are out, according to the survey we did anyway. Even the ones that are quite

enthusiastic are still very much in an experimental stage, but I think that's going to transition very quickly into much more meaningful activities.

The “experimental stage” to which Beckett refers is very much a reality for news organizations who are cautious to fully implement generative AI into their organization for a variety of reasons. Jay Allred explains why his newsrooms have yet to completely embrace fully generative AI into their story-writing practice:

We have a few newsrooms that are starting to get to that point where they're like, 'well, we like what we're getting, but I think let's try generative, and see what it does to create some variants and stylistic changes and those kinds of things. But we don't want to automatically publish, we want to publish a draft.' No one that I know of is using generative, fully automatically, to where it's literally writing the stories. I think they're gunshy, frankly... 'oh maybe I don't want to be mobbed by people on Twitter because of this.'

Allred explains how newsrooms are “gunshy” as a result of the potential backlash from the public. Whether it’s getting “mobbed by people on Twitter” or excoriated by fellow journalists, news outlets certainly want to avoid any damage to their brand that premature adoption may cause. Ryan Restivo has had direct experience with newsroom resistance to his app, YESEO:

They all just don't want to be the next “that one.” That's the fear, right? The look from the mindset of we don't want to be the next “these guys.” With our headline with AI next to it “AI gone wrong: [our outlet]?”

Restivo underscores the idea that newsrooms don't want to be the one that makes the next mistake. In order to overcome this resistance, Restivo explains how, in his experience, he can help demystify the AI tool adoption process:

...how do you get through those training boundaries and the resistance in newsrooms? You know, for me, it's answering people's questions directly, being direct with them. I've met with newsrooms that have wanted to use it or use it more often. And in being open and just saying, hey, well, what are your questions right there when I go to conferences, you know, there have been a couple of conferences where it presented a Q&A portion, right, I'm the best person who knows about my thing better than anybody else. So they can ask anything they want.

Ultimately, a common thread among respondents was the importance of newsroom proactivity with AI. Both experts and those involved in newsrooms (as employees, partners, sellers, etc.) recognize the importance of a baseline level of strategizing and planning around AI tools. Independent of excessive caution or recklessness with adopting these innovations, an organizational strategy on how to approach AI—and generative AI more specifically—is essential, as explained by Jared Schroeder:

I don't think there's a lot of benefit in news organizations saying, 'we're just not going to do anything about this.' I think that's the only sure way to fail. I think it's better to say 'how are we going to approach this thing?'

5.3 Business Models

Another common area of interest for respondents was business models: namely, how artificial intelligence fits into the economic strategies of news organizations. Individuals involved in the research and/or production of automated news content articulated the impact that AI technology has had on traditional newsroom business models: the organizations themselves, their processes, and the legal and competitive implications of AI in news production. Automated content creation is just one of a multitude of areas cited by respondents as a domain of transformation by artificial intelligence.

The systemic-level impact of AI, mentioned by Charlie Beckett earlier, impacts not just the bottom line of news organizations, but also the nature of the content that is produced. Chris Dinn, founder and publisher of Canadian news startup Torontoverse, articulates what he feels is the systemic impact of automated content production based on his experience with working with the technology:

...my feelings on this have evolved over time, as I've spent more time working with the technology directly. I'm much more, I think, you know, I think that AI is going to do a lot more content creation here than we realize. But I think it's not going to result in less, it's going to result in better content and more content, hopefully, I think it's, if we're lucky, it's going to result in a revived news industry, on the grounds of fewer people being able to accomplish a lot more.

Dinn's observation of a "revived news industry" based on "better content and more content" is an optimistic viewpoint, but is informed by his own experiences with AI easing business processes. He continues:

...one of the things that news really struggles with is that it's hard to monetize that volume. It's hard to produce enough content, to pay for the staff that it takes to produce good content, I think that's gonna get a lot easier with AI... I use embeddings a lot. So pulling out embeddings, working with LLMs, directly working with related models, image generation models, and vision models, which, you know, are LLM-adjacent, is the relevance.

Much like how AI as a tool at the individual level is primarily adding value and impact through processes adjacent to text production, so too is the value at the systemic, business model-level. Dinn's description of "LLM-adjacent" models in the embedding, image generation, and vision spaces reinforces the idea that AI will help news produce enough content to pay its bills and pay its staff. However, not all respondents shared Dinn's optimism in the future potential of producing content at scale from a business model perspective. Charlie Beckett adopts a more conservative perspective:

On scale, I would say scale is over, and has been for some time. Frankly, over the last five years or so, especially with the decline of the utility of social media for journalism, it's just not a very good place to market your content anymore. It doesn't like us anymore. And the shift towards a subscriber model. Now, obviously, you want scale in terms of subscribers, you want as many of them as you can get, but it may be that you survive with just 100, say you're a tiny little niche organization, you just need 100 subscribers, that's it. You don't want 1000, you don't have to make an effort to get 1000, if that's sustainable for you.

By claiming that “scale is dead,” Beckett counters the traditional media focus on scale, noting that the large-scale distribution model facilitated by social media is losing its effectiveness for journalism. This reflects a shift in the landscape where traditional methods of content distribution are becoming less viable. Beckett continues by explaining that not only does the decline of social media utility contribute to the death of scale, the sheer amount of journalistic content available leads to a supply and demand problem as well:

And I also think that scale is dead in the sense that the world does not need more journalism, there's too much journalism out there. The brand that you consume will create far more content than you will ever be able to consume in a day. So rather than scale per se, for some organizations like The Mail online that survive by pumping out loads and loads and loads of content, and there are always going to be content farms.

Now that the amount of content far surpasses what the world needs, and there is “too much journalism out there,” news organizations must tackle the problem of looking past the idea of content when considering their purpose in the future. Based on his experience as a journalist and historian, Jeff Jarvis also recommends looking past the idea of content. His observations on newsroom infatuation with content, monetization, and commodification reflect his opinion that such a focus is a byproduct of the “Gutenberg era” of print, which is coming to an end:

...the Gutenberg era created this conceit of content, that what we produce is content that fills things that has a beginning and an end and Alpha and Omega, it is contained. But the problem is that it seduced us into thinking that our value was resident in this thing called content and that

we create content and people should pay for it. That's not where the value of journalism lies. The value of journalism is as a service to communities and their discourse and to democracy and society. Content is just one tool we have. And so if we imbue content with all of our value, and then content becomes commodified, we're screwed.

Jarvis presents a dichotomy here between what he calls the “conceit of content” and journalism “as a service”. The belief that the purpose of journalism as a business is to manufacture content has motivated the organizations with which Jarvis has worked to fight to protect their content from what they view as theft using copyright as a tool, as opposed to being concerned with outcomes instead of products. Jarvis, as a journalistic and content creator, observes that in a world in which information has an infinitesimally small half-life, content loses its value. And if newsrooms have binded themselves to content as the centerpiece of their business model, Jarvis opines that they are “screwed”.

Beckett sees the future of journalistic practice being more about the optimization of content than content at scale. He describes the promising possibilities of personalization that has been borne out in his own observations and experience:

...the real success is going to be about optimization of content. So it's going to be what do you want to read today? And how would you like it? Would you like it in short bullet points? Would you like some audio? Would you like a bit of video to go with that one, you know, etc, etc. And when do you want it? In the evening or in the morning, you know, and what kind of topics and here's some surprising topics you haven't thought about that you might be interested in as well. You know, more like a restaurant, frankly, than a fact. And that's not about scale.

Beckett's idea of content optimization being "like a restaurant" emphasizes the personal aspect of the future of content. Other respondents echoed this idea, not just in terms of content being personalized for users by other writers, but also individuals creating content for themselves. The computer scientist and generative AI startup founder in the U.S. mentioned this idea:

...another reasoning that I heard... is this idea that everybody will be able to create their own stories, and you generate your own content. And so the overlap between people's knowledge bases is going to almost completely disappear. Whereas right now, despite all the many sources, we still have some commonality in our content that we consume? And that's really scary to consider, would you just generate content specifically for you?

What she describes as the "overlap between people's knowledge bases" is the value at stake in an era of highly personalized content. Furthermore, if this highly personalized content is able to be generated using artificial intelligence, the current information glut could get even worse. Jared Schroeder describes what he believes to be an upcoming struggle that human journalism will have in a world of AI:

The network era created an information universe where information was no longer scarce. It is the opposite of scarce—it is everywhere. What's going to happen when AI is generating more content than humans are? By exponential levels—by 100 times or 1000 times? How is journalism with humans going to compete in an era where the internet is basically full of just AI generated information?

Schroeder here considers how human journalism will “compete” in a landscape where the authenticity and origin of content may be indistinguishable, challenging the traditional value and trust placed in human-generated news. Another important dynamic in the competition between human-generated content and AI-generated content is the issue of copyright: large language models training on works published by humans. Beyond considerations specific to the model training process, many respondents agreed that AI training on published news articles poses a threat to traditional journalistic business models. Schroeder describes his own experience with an AI news report tool, Perplexity, and the challenges it poses:

[Perplexity is] kinda like a search engine based on ChatGPTs undergirding programming. And ...Perplexity basically replaces a news report... [i]t'll cite everything, and it'll cite traditional news sources... [and] so people won't look at journalism anymore. Because they'll just go to the AI tool and have it tell them what's going on. And it's a simple step to say, Well, maybe you don't go to Biloxi and ask, you just get a newsletter from a company that provides you information that's from other news sources and they never pay them.

Schroeder also opined about the OpenAI-New York Times lawsuit, which at the time was in its early stages, and its implications for the journalistic business model:

...a separate thing to pick up and think about is what has happened to journalism's financial models in the last 25 years. They've gone from good to gone. So one of the things that we have to keep in mind is whatever is decided in this case could have a lot to say about the future of

journalism, like: should we have it? Because if it is ruled that it's not a fair use for these companies to train these tools without any kind of permission or paying any amount of money for the use of this work, then journalism just got a new lifeline because it has a revenue stream that it didn't have before.

Schroeder is referring to the idea—relevant in the lawsuit—that OpenAI should be licensing content from publishers rather than scraping it from the web. Striking partnerships with publishers can allow for AI companies to avoid issues related to copyright lawsuits. Charlie Beckett discussed the example of the Axel Springer-OpenAI partnership:

...the Springer deal with OpenAI just showed us how rapidly this has become a substantive business model issue, as well as a practical thing about how journalism is produced and consumed. And it's a kind of landmark deal, because it's not just a payoff, it's not kind of 'we want some of your money'. This is a news organization saying, 'Okay, you can use our data to help train your wonderful models,' and that's just really, really seismic, potentially.

In terms of business models, these partnerships have mutual benefits for both sides. Some respondents felt that these partnerships will become commonplace as a means of circumventing legal battles and forging sustainable long-term economic deals. Schroeder explains that these partnerships would be beneficial for customers and subscribers as well:

...what if you create New York Times GPT? Lots of organizations are doing this, researchers are doing this, creating their own specific tools, AI tools for their research or their news, where it

becomes a benefit for customers that are subscribers... they can use this tool to search the entirety of, they can basically make their own news story. If instead of going to Wikipedia, you go to New York Times GPT.

Ultimately, as respondents noted, future business models for newsrooms may need to focus less on protecting their content from competitors (whether they are fellow news organizations or fellow information providers such as OpenAI). Instead, we may see news outlets entering into agreements with these competitors in order to provide the most desirable services possible for customers, subscribers, and the public at large to sustain profitable business models.

5.4 Audience

Another common theme among respondent discussions of their experience was their relationship to an audience. Journalists and news producers articulated the importance of audiences—news readers, product users, and their communities—in the ethos of journalism, and how increased reliance on artificial intelligence affects the role of trust in the source-audience relationship. Interview respondents described their audience differently based on their professional experiences, with some emphasizing the role of audience as community, and others emphasizing the role of audience as market.

A local news project leader shared with me her experience as a practicing journalist. She emphasized the relationship between journalists and local communities, which she described as difficult in a variety of ways:

And through the course of that work, I got frustrated with... sort of the disconnect between the journalism and the community, in many different constituencies, and feeling that there were some issues there. Also, feeling encouraged by some things that we were able to experiment with in terms of doing things that we might now call solutions journalism, or engaged journalism, both in the US context and internationally.

She describes feeling both “frustrated” and “encouraged” with approaching how to engage these different constituencies. Furthermore, she goes on to articulate two major needs in this engagement: assessing the communicative health of a community—or public sphere—and the connections within the journalistic network:

...there's some resonance to public sphere concepts in that it's all about the looking at local storytelling networks and looking at the connection between these different parts of the storytelling network, which they conceive of as community organizations, residents—networks of residents, and local community—journalists, local journalism...

She goes on to articulate, from her experience, the benefits she observed of synergies between these three parts of the “storytelling network”: community organizations, residents, and local journalists:

...in neighborhoods where they found strong connections between those three parties, there tends to be a shared sense of belonging, a shared understanding of what the local issues are. And, you know, with the potential of having a shared conversation around how to address challenges in

the community... And so that's sort of where I go when I'm looking at different communities and trying to understand: what is the communication health of a community?

The important takeaway here is that these connections between organizations, residents, and local journalists are important at a baseline level in order to have a “shared conversation” about the challenges, let alone address them. For this respondent, at the local journalist level, this starts with accurate reporting that is informed by context and history:

I think for a lot of the work I do, there's a question of understanding the context and history of a place and a community. And a sense that you don't do a story about gentrification without talking about redlining. Don't do a story about how crappy things are in my neighborhood without explaining why and how they're tied to a longer history of all the various ways white supremacy has affected their community.

The extent to which AI can assist with this process determines its effectiveness in strengthening the organization–resident–journalist network. Many respondents aligned with the sentiment that in order to integrate technology into the journalistic process, it requires not only understanding the audience of the journalistic product, but also the technology itself. On bridging the gap between newsrooms and local communities, media expert Damon Kiesow explains:

it [bridging gap between newsroom and local community] really is going to require a significant reorganization of how we think about technology, most news organizations, even some of the many of the newer digital native sort of orgs really still look at technology as a neutral tool.

Kiesow elucidates the idea that generative AI is not merely a “neutral tool” but is shaped by and shapes social contexts and organizational structures. This concept aligns with the idea that the development and implementation of technology are heavily influenced by social, cultural, and economic forces. Kiesow sees newsrooms accepting new technologies deterministically, when there should be a more focus on community impact:

...generally, when humans run into a problem, they can't understand the system. They either reject it entirely, or they accept it unquestioningly. And we typically tend to just accept it unquestioningly, based on what the vendor told us. And now we're accepting the vendor, the vendor's logics, that technology's logics, without having an understanding how that affects our, our business, our newsroom culture, our community.

The importance of cultural considerations was echoed by Gregory Gondwe, a professor and researcher. He illustrates this point by explaining the problem of designing a tool with Western culture in mind and deploying it in a non-Western context:

And also when let's say they want to include a culture within Africa, they perceive it in terms of Western culture. And that's where they are building the idea of civility and incivility comes in, so what is civil is defined by what the West think, is simple. So, when these generative technologies come and they design something based on the concept and on the language and the culture, on the Western culture, ignoring the baseline of the cultural values within the context...

“They” in this context refers to tech companies that market their generative AI products outside of the United States. The key question here is tailoring the design of the technology to how the audience—in this case, the community and its cultural consideration—receives the technology that it is adopting:

...so, of course, questions of hate and all kinds of stuff are not cultural things, but there are certain things that might sound civil in some cultures, and yet they are not similar in other cultures, and those are not necessarily considered.

Beyond the audience as a community, many respondents described their audience as a market—describing them in terms of users, addressable revenue opportunities, and discerning consumers—and emphasizing the process of interaction. Ryan Restivo articulates the challenges of audience integration and applying a small tool such as YESEO to a large audience:

...when you take a small tool, and you try to unleash it on a large mass audience in a certain amount, and within a certain way, it's going to be very difficult, right? Over 300 workspaces have installed [YESEO] right, that means they have a chance to use it to improve their content. But at the same time, they still have to make those actions.

Restivo’s audience of users is different from the audience of a traditional news organization—his users are newsroom workers themselves—but there is still a source-audience relationship. He describes the feedback loop between himself and those who use his product:

...the first month of this, I got good advice: just watch how people use it, right? You'll see in the data, how people will use your thing. And when it goes out in the wild, and you have all of these people using it, you get a kind of good sense of just how, like how somebody else might use this thing? Or what adjustments need to be made off of it? How are they using it wrong, that you need to now change? Or show somebody, Hey, use it this way.

For Jay Allred, his audience includes both readers and newsrooms. When describing how he launched his company, he articulates the challenge he faced in terms of existing solutions to audience demand:

And we knew from our experience at running a local news organization, we knew that our readers... were interested in high school sports. And we knew that there was no way we could cover every game. And particularly with prep sports... it ends late in the evening. You've got people trying to run down scores on social media and Twitter to try to figure out what you know... it just was a difficult thing to do.

Based on what his audience was interested in, as well as his own perceptions of what was possible on the product side, Jay explains how he arrived at the idea for his news project:

And so we found this combination of an approachable dataset, an unknown addressable market, which are people that are interested in high school sports results, but aren't getting them and an experience and how to write those things? Or like, how might you write the templates for those that just kind of lead itself into, okay, we have the right pieces to make a project out of this.

For Jay, conceiving of an audience (or potential audience) as an “addressable market” was a key factor—one of the “right pieces” to the puzzle of how to solve the problem of satisfying demand for prep football game recaps. Chris Dinn also emphasizes the importance of thinking about his audience in terms of “economic potential”. As the founder and publisher of Torontoverse, he has a vested interest in the economic viability of his publication and keeping his audience happy:

So how can we as people who want to publish content, that want to make news, that want to keep people informed... how do we capture the economic potential of our audience and make it available to businesses or partners that might help fund our journalism? I really feel like that's at the heart of what we're doing.

He goes on to contend that AI facilitates the achievement of both of these goals:

I think LLMs offer an incredible opportunity to do this. Imagine targeted advertising that people don't hate, because it's able to leverage your market and also a deep relationship with the customer. I think there's a huge amount of opportunity here. But I think we're at the very, very early stages, seeing what that might look like.

His contention that “we’re at the very, very early stages” implies that the future holds potential for AI to unlock the prospective economic value of audiences for publishers at an unprecedented level. However, Dinn and others note that there will be drawbacks to AI displacing search engines and other common channels:

...the main thing is consumers are going to change their behavior based on this. And that's going to mean big changes for publishers, which goes back to... [it will be] harder to find an audience.

Not only will consumers “change their behavior” in terms of their information seeking patterns, in the case of articles that are produced by or with the help of AI models, they will also develop awareness of AI authorship, either through bylines or skillful discernment. Jared Schroeder recalls his realization that audiences are becoming aware of such phenomena:

I was reading something yesterday... about the fastest goal ever scored in the European Cup. And this guy from Austria scored in six seconds. And I read the report thinking ‘this is so interesting.’ At the bottom, the human comment, the first one was ‘this is totally an AI story.’ I was like, ‘Hey, I never even thought that,’ but whether they're right or not, it shows the audience is already thinking about these things.

5.5 Trust and Transparency

Another common experience among respondents was how, based on their experiences, they felt that AI affects the role of trust in the source-audience relationship. The U.S. east coast data journalist articulates his thoughts on AI in data journalism specifically—more so in the reporting assistance process rather than in a text production capacity:

So hopefully, people trust articles where the reporting process used an AI component or data journalism example, for the same reason, hopefully, that they trust...that we care really deeply

about getting stuff right. We lose sleep when we get stuff wrong. We have processes that protect against that. And so I think they should, they should trust us, but not because of the AI.

He mentions that “they should trust us, but not because of the AI” underscores the importance of trust as a general journalistic value. If a reporting process were to utilize AI tools or techniques, that usage in and of itself should not affect trust. He goes on to explain that the role of AI in gathering information is similar to that of a human, and as such, its role need not be disclosed any more than that of a human serving such a function:

And I think there's a huge number of cases in journalism where this methodology is useful. Where the AI is serving the same role as the tipster who sends you an envelope with no return address that says, look into these three things. Once you look at the three things, that's your story, the fact that it showed up in an envelope doesn't matter. Likewise, when the AI highlights these examples for you, you investigate them with normal journalism skills and write about what you found.

His evocation of the role of AI as “tipster” underscores the idea that when AI is involved in production in a more indirect manner (such as gathering information and otherwise assisting in the pre-writing process), it is not supplanting the role of the human author:

And the fact that you used AI is like a fun story to tell at the bar to your friends, and not something you really need to explain to readers that much. Because it isn't what they need to know. It just isn't the story.

While he contends that AI involvement in these procedures “isn’t what they [readers] need to know,” other respondents contend that this lack of disclosure reflects a larger issue of transparency in the journalistic process beyond AI. Joe Amditis articulates this idea:

I think journalism as an industry also has a lot bigger problems to grapple with when it comes to transparency and revealing the sausage making process. How many editors and newsroom employees other than the original reporter or op-ed writer, touch a story before it goes to print? I don't see any editorial bylines, I don't see the name of the person who wrote the headline, I don't see that kind of transparency.

Amditis sees the issue of “sausage making” as one in which readers are unaware of the multiple entities who “touch” or influence a story before it is published. The increased awareness of AI in the news production process offers news organizations and their readers, according to Amditis, an opportunity to have a conversation about journalistic black-box processes more generally:

I think there's a larger conversation with regard to the trust and transparency that newsrooms operate on among their communities, that includes AI, but it also speaks to a larger issue of revealing the curtain and pulling it back.

For Amditis, this “conversation” is important to have because of the different expectations traditionally held for different “kind[s] of transparency”. David Cohn identifies the question of

what kinds of transparency are required for journalism as an unresolved and important issue given the presence and varying magnitude of AI contributions:

And the question is: transparent about what? Because let's just say I write a story, and I entirely write it myself, but I use an AI to help me come up with a headline. In fact, I say, here's the body of the story. Give me three potential headlines. One of them I thought knocked it out the park and that's the headline I go with. I would argue that I don't need to disclose that headline. Because, again, the headline is wholly from the body of work, which again, I'm saying is wholly written by me. So the question is: what are the material or substantive contributions from AI that merit disclosure?

The extent to which AI is influencing the production of journalistic output differs from report to report, from article to article, and from organization to organization. Similarly, different organizations are dealing with Cohn's question in different ways. For Jay Allred and his high school sports automation project, the answer lies in the deployment:

So the focus that we're having on this is we are all about effective deployment. So like, all we think about is like, what does an ethical deployment look like? What is an effective deployment? How does deployment affect the communities that are there? What do you do with trust? What's transparency? How do you deal with all of those kinds of things?

Allred explains that for his organization and the newsroom customers they serve, transparency is about being clear with the audience about what is happening, why it happened, and what

happened—this applies not just to the content of the story itself but also the use of AI is producing that story:

So for every AI story that we write and publish, there's an author associated with it called Auto Newsdesk. And that author tells the [reader] 'I'm a helpful AI assistant that gathers and writes short articles about high school sports results' ...we're clear, we work to be very clear, and we advise our customers to be super clear with their readership around what they're doing.

Not all news organizations have followed these principles accurately, however. Respondents discussing transparency were quick to note examples of outlets experimenting with the technology without informing their audience, such as Sports Illustrated, CNET, and Gannett. Jeff Jarvis observes that this type of experimentation can reflect poorly on the reputation of these outlets:

I think we see stupidity going on. We see some smart people going on, obviously. What GeoMedia and CNET and Sports Illustrated and Gannett did was all pretty stupid.

Given the ambiguity of the question Cohn poses, many of these organizations experimenting with AI-powered news production who fail to disclose the aspects of their process that “merit disclosure” have learned the hard way and suffered reputational damage. In the case of Sports Illustrated, who employed a company that produced articles for its website written under the byline of authors who apparently did not exist, their failure to disclose their decisions, more so

than the decisions themselves, was the central reason for the public backlash. David Cohn explains:

...what Sports Illustrated did, it wasn't so much the crime but the cover up. I think the idea that you would not disclose it and pretend as though these are people is, I think it's worse than if they had just done this and said, like, 'AI [did this]...'

The notion that, as Jeff Jarvis says, there are both “smart” and “stupid” ways to experiment with AI aligns with Jared Schroeder’s experience as well. He offers an explanation for why examples of both failure and success will manifest themselves in this era of experimentation:

I don't think we have a list of best practices yet. I think there's going to be a lot of experimenting, there's going to be some failure and embarrassment, Sports Illustrated... there's going to be some missteps, and we're going to learn from them. There's probably also going to be some shining examples.

The lack of “best practices” surrounding disclosure of AI usage in news production is a factor that has enabled both successes and failures in these early stages of generative AI journalism. Ultimately, as respondents have alluded, trust in written work requires an ongoing negotiation between the reader and writer. Joe Amditis asserts that in his experience, in order to manage a relationship of sufficient trust with its readers, a news organization needs to have a modicum of self-awareness:

*...you have to have an understanding, at least at some level, of how you are viewed as an organization as an industry by the people that you're purporting to serve. One of the big problems with journalism, as an industry, especially some of the more traditional and large legacy organizations, is they have no self awareness, they don't have a sense of humor; they don't have an ability to look and reflect on, you know, the kind of s*** that they've done over the last couple of decades or so... so this is an opportunity to open those doors and pull back the curtain. But it has to be done with a sense of history and understanding of their own role in that.*

Amditis and other respondents have learned from their experiences that if news organizations are consistently engaging with—and being transparent with—those whose trust they seek to earn and maintain, then they will be more likely to succeed in their experiments with new technologies such as generative artificial intelligence.

5.6 Guardrails

Most respondents discussed their experiences and knowledge of the risks surrounding the integration of AI in journalism. They emphasized the importance of implementing organizational safety measures to responsibly integrate artificial intelligence into the journalistic process, minimizing the potential for harm caused by the technology such as misinformation, hallucinations, and biased training data. A common thread among these responses was the term “guardrails,” a word which literally denotes a safety barrier designed to protect against some sort of danger. In the context of AI, the word has been adopted in a similar context to describe measures that mitigate risks. As Jeff Jarvis explains, developing such measures is a challenging task:

...it's a general machine. This is why I think guardrails are basically impossible, because it's a general machine, and there's no way to predict everything bad that anyone could ever decide to try to make it do. That generality is what excites the AI boys a little bit too much. But it also makes it, I think, very difficult.

One of the risks most commonly discussed among respondents was misinformation. The east coast data journalists summed up the relationship between AI and misinformation succinctly:

...these technologies make it really easy to lie at scale.

At the same time, he clarifies that this ease of producing misleading content does not necessarily translate into increased risk for readers:

I'm not sure it makes their lives all that hard, much harder, because lying at scale is already really easy. There's a lot of it. And the solution still is kind of like finding people who have the right process.

Charlie Beckett echoes this sentiment, contending that the worries about tech such as AI amplifying misinformation may be overstating the issue, or at least misjudging the fundamental root of the problem:

I think we can over-exaggerate this. In many ways. Yes, there will be competition between I mean, the bad people are already using generative AI, of course we know that. But disinformation is a human problem. It's not a tech problem. And it doesn't have a tech solution. It's a human problem, because it's governments like the Russians, or its people trying to make money. Or it's conspiracy theorists. And it's us sharing the stuff and reading the stuff, you know, that means this exists. It doesn't exist because of the tech. I mean, the tech obviously enables it and amplifies it, and so on.

He also calls into question the risk it poses for readers:

There's also a big question about how effective it really is, you know, a lot of research is telling us that a lot of the disinfo has very little impact, in terms of changing people's minds, there's a kind of general impact, which is that it makes people more cautious and cynical about trusting things. But that's possibly not a bad idea. You know, it could be that disinfo is actually quite a good media literacy exercise.

David Cohn asserts that the risk level of misinformation is different for journalists versus the general public:

...in terms of risks for journalists, that [misinfo and disinfo] is the one that probably is the most potent. And funny enough, we're not concerned that we're going to be producing it, it's more that that is going to run rampant, and somehow there's going to be this glut of content.

He mentions how as journalists, there is not a concern that the misinformation is coming from oneself, but rather that the “glut of content” is emerging from elsewhere. This lack of concern downplays what Cohn regards as the more “potent” threat of misinformation compared to the risk it poses to the public:

And in the public and the general public, I like to think that that's something people are aware of... I'm not speaking on behalf of them, obviously. But I think the public is aware of it, but I don't think they're quite aware of how acute or pernicious it can be.

Developing guardrails for misinformation risk requires first and foremost an understanding of the problem at its root. Beckett explains that solving misinformation requires regulatory measures:

So I think we have to understand that we need to figure out how to try and counter it, we need to understand the sources of the disinformation, and we need to understand why people are doing it, and what the real impacts are. And that, in a sense, is a regulatory issue.

Based on his experience, Beckett explains that some news organizations aim to solve the problem of misinformation with AI—fighting fire with fire in a sense—which would necessitate the implementation of what he calls “structural mechanisms”:

...there's real structural issues around if you want to use AI to solve the problem, then you're going to have to have some sort of structural mechanisms built in such as watermarking generative AI content, for example... that's a big ask. Some people will do it, because they want

to signal their virtue. So I don't know, if you've got a large language model that's producing stuff, and you want to be trusted perhaps you should say we're going to whatever... mark or watermark, anything that's produced using our tools, so that other people can spot it. You know, that would be nice if they did that. I don't hold out much hope...

Beckett's description of watermarking as a "big ask" echoes the prevailing notions of the technique as promising but, unfortunately, unlikely to work due to a variety of reasons. Indeed, guardrail implementation for disinformation is not an easy problem, and for respondents who spoke about their experiences, they tended to agree.

A related issue that some respondents brought up from their experiences with generative AI content was the concept of hallucinations. Similar to the use of AI tools to deliberately produce false or misleading content, hallucinations result in inaccurate or nonsensical outputs. However, as Jeff Jarvis notes, the term "hallucination" may not even be appropriate as it implies a particular relationship between the language model and truth:

I object to the notion of hallucination or even the notion of lying, because that would assume that the LLM has a sense of truth and it does not. So we need different nomenclature for this. I think I think that's important so people understand that.

Similarly, Charlie Beckett asserts that the effectiveness of large language models should be seen as tied to truth:

...it's not a truth machine. It's a language machine. All those things... it doesn't know anything, it has no intelligence, it doesn't have any sources, and therefore, we're all going to find out, as with every sector, every sector, from law to retail, et cetera, is going to find out well, how effective is it, really?

While it “has no intelligence,” an LLM can generate accurate information from its training data. As Jarvis explains, restricting an LLM to a corpus works better in terms of predictive text and accountability:

...if you restrict it to a corpus, number one is that it knows what it's calling on, it must call on that corpus of data using the skills again, to know how to put one word after another and to predict that word, but it does so from that set of data. Number two, is that it has accountability. What I've seen those that are trying to do this, is that it... cites sources. Now it's got to understand — I use the word 'understand' guardedly, right—but it has to recognize a claim and the source and the different citation. But the more that happens, the more can be checked, the better.

Jarvis mentions the fact that a model trained on a restricted corpus “has accountability” in the sense that it can be “checked” or audited. Auditing is a potential guardrail for hallucinations with which some respondents have personal experience. Chris Dinn explains how, as publisher of *Torontoverse*, he audits his AI agent tools for hallucinations:

We have [AI agents] going on downloading websites and coming back and assembling data, facts, and citations together. We have vector database lookups with citations in them, etc. One of the things you spend a lot of time doing is looking at the citations, looking at the facts that it listed and trying to figure out if, actually, the facts are true or not. And then seeing if it was hallucinating, if it was hallucinating or did it pull it out of its weights, were they accurate, and really chasing all of this stuff down.

His AI agents assemble data, including citations, from websites that have been downloaded into its corpus. Because the output of the AI agent is critical for the end users – the newsrooms that use these AI agents that Dinn develops – accuracy is vital, and as such, guardrails are essential. This is, as other respondents have noted, an issue of accuracy rather than truth—making sure that the output of the AI model is an appropriate probabilistic outcome based on its training data.

Dinn continues:

...we are building this product — the Torontoverse CMS interface — to provide accurate information for people to write accurate articles. You can't just assume that it's plausible enough sounding—and it often is plausible enough sounding—you'd really dig and see, is it true? If you do things like ask for citations, that increases the accuracy dramatically; it has to produce a citation for everything. But there are still times where it says things like: is this true? Or is it not true? I think this sense of figuring out whether or not what you're looking at is a hallucination or an actual fact has become a lot more important.

A language model such as the one powering one of Dinn’s AI agents is only as accurate as its training data, and as such, biases that are present in the training data are often reproduced in the output. As many respondents observe, the issue of bias in training data is another reason why guardrails are so important in AI development. The U.S.-based computer scientist and startup founder articulates this issue in more detail:

...machine learning is fundamentally biased towards the data that is trained, it’s not even biased on it. It’s just entirely dependent on it in every conceivable way. But when you have massive models trying to interpret data, it’s going to talk like people do.

When large language models are trained on enormous amounts of human data in order to “talk like people do,” decisions about what gets included and what gets left out of the corpus are made that influence the level of bias engrained in the model. Some respondents spoke to the issue of unrepresentative data sets being used to train journalistic AI tools. Gregory Gondwe explains how, in sub-Saharan Africa, journalists utilizing ChatGPT had to deal with model outputs that did not take local context into account:

...most of these technologies are not designed for the local people. They’re designed for the elite for lack of better words. Of course, the push from many people has encouraged these organizations to start considering the local people. Now, for example, ChatGPT has languages like Swahili, which is great. But even at that time, I think most of the journalists actually bemoaned the fact that it wasn’t that representative. The data sets were not that representative in the sense that we didn’t see the common language being used within the context.

Inclusion of diverse data in a training data set, as Gondwe explains, is not as simple as including a sufficient amount of data such that it can produce Swahilian text; there must be enough data for ChatGPT to be able to recognize context clues embedded in the patterns of data on which it trains. He continues:

So maybe it wasn't informed enough to explain what was happening within the context. And for lack of a better word, they thought 'yes, we have been given this, they have included our language, but they have squeezed out the juice'.

Gondwe describes how the decision-makers training large language models “squeezed out the juice” by failing to construct a data set that was sufficiently representative for use cases in sub-Saharan Africa in which journalists relied on it for research.

Respondents spoke to another issue inherent to language models trained on Internet data, which is the more broad issue that text scraped from the Internet is not necessarily representative of human language. Chris Dinn brings up the fact that models also need to be updated in order to account for new data, and models trained on Internet data with an old cutoff date are unrepresentative:

...we have to come to some acceptance of you don't want AI to ossify culture a little bit or ways of thinking by making the ubiquitous views of the internet from 2004 to today, the Common Crawl data set people are using to train or pre-train their models comes from... you don't want

that to be forever the way that society operates, because... forever 90% of the content we generate was based on this subject matter base and doesn't change very much.

He continues by recounting his own experiences developing technical solutions to address potential bias in the training data on which his AI agents at Torontoverse are trained:

And then to try to solve the problem itself with technical solutions that mitigate, you know, that take the sharp edges off the LLM, and make it less risky, create content with that is providing more accurate facts.

Dinn's articulation of debiasing as taking "the sharp edges off" evokes the idea of smoothing, which in data science refers to a process of eliminating outliers from datasets to make the patterns more noticeable. He goes on to describe this process of risk mitigation in more detail:

...grounding the generative prompts very well, so that building a database full of data from the city of Toronto website, doing a vector search before you generate or try to answer a fact or bring a fact into the prompt history, digging out the necessary embeddings, including them when you're trying to generate a piece of information that you're gonna use, or when you're gonna review, a bunch of other information can help a lot.

Dinn mentions several important steps here: prompt grounding, vector searches, and "digging out embeddings" – in other words, extracting the mathematical representation of the underlying

relationships and patterns between points in the training data. He acknowledges the difficulties of this debiasing process while still remaining confident in the capabilities to mitigate risk:

But it's an art and... six months ago, I was optimistic it would be very easy to do. Six months later, I still believe we can do it. It's just a harder and more nuanced problem than I thought.

The challenges of addressing biases in training data are not just limited to their nuance. Dinn observes that what it means to “eliminate bias” now may change over time, and quicker than we anticipate as a society:

...biases are a moving target, like, I think we might say, ‘today, let's eliminate bias’, you know, and what our vision of that is versus what the vision of that is 25 years from now might be very different. The ideas of bias, what groups we think are important, how people self-identify might be different in 25 years, and certainly with the coming change of AI, maybe... it'll be even different. Maybe there'll be people that identify as pro-AI and anti-AI or whatever. There could be things that we consider important that we can't anticipate.

He explains that, in the effort to include the widest range of perspectives and experiences in a particular set of data, practitioners must acknowledge that these perspectives and experiences will continue to evolve with society, and as such, inclusion will be an ongoing need, in training data and beyond.

5.7 The Human Element

Researchers and practitioners on the media, business, and technology fronts all expressed the importance of having a “human in the loop” in order to harness the innovative potential of AI technologies. The constantly evolving dynamic between human and machine from a labor perspective was also a prominent theme among those interviewed. In particular, respondents with machine learning experience noted the value of humans aiding computers in making the correct decisions in building large language models and selecting the most critical data needed to refine them. Outside of the realm of generative AI, respondents emphasized the need for human initiative and oversight at a procedural level. The U.S. east coast data journalist explains his process:

I use AI sort of as part of the exploratory process; we ended up using a much simpler methodology. But the use of AI was just the way to get from A to B. To get from 'here's the data we have' to 'here's the answer to the question that we need'.

He makes an interesting point here about using AI “to get from A to B” which underscores the role of machine learning processes for a data journalist, which we have already seen is a tool with multifaceted uses. Ryan Restivo explains why his tool, YESEO, inherently requires the element of human decision-making:

So the idea being that because it's in Slack, you can't just hit a button that's going to pump it into your content management system, you have to make a decision. So at some point, a human—and it's probably one thing you hear a lot—a human's always going to be in the loop. And they're the

one that's gonna have to decide which pieces of this are good, which are not, and how can I best move it forward?

Restivo's use of the word "decision" is important here, in the sense that there is a chain of choices that must be made in order for a headline generated by the YESEO tool to ultimately be published. These choices are made in part by AI and in part by journalists. Ultimately, the human journalist is the one that must "decide" on whether or not the quality of a particular AI output meets the (again, human-decided) standard for publishing. Restivo further elucidates that this joint process is what the norm should be in terms of AI suggesting aspects of stories to journalists:

I believe that a human has to be in the loop on all this stuff. Like with data, we're trying to get people better data to make these decisions, people are still gonna have to do the work to get this right. If a person's still gonna have to write that story to then get the suggestions to then figure out what they want to do with that headline, people are always going to need to be in that spot to make decisions.

He compares the need for humans to make decisions in journalism to similar needs in other industries, illustrating the instrumental role of humans in hybridized processes across a multitude of domains:

...when you look at, like, instant replay reviews and like football or baseball, the person's always going to have to decide something. And you still have to challenge or do something to make it

work. So we don't always want to just be automatic and change, it's that we want people to get towards their better decisions and people to make the decisions and get towards those outcomes.

In Jay Allred's newsrooms, his journalists use tools similar to YESEO to optimize content production:

...we're using tools now to help SEO optimize the content that we're producing. So, 'give me better tags, give me some headline optimizations.' Okay, so one of two things would happen in that scenario: either a) it wouldn't get done at all. But there would be no SEO optimization that happened at all, or b) it would have been done, but it would have taken longer.

In this case, the AI is enabling an optimization process that wouldn't take place at all otherwise. However, he explains that these technologies help with the speed, quality, or sometimes both, of reporting:

So broadly, these tools are speeding the process, or if they're not speeding the process, they're upping the quality. And sometimes they do both, sometimes it's a trade off, it's one or the other.

Ultimately, beyond the enhancement of speed or quality of the output itself, Allred espouses the goal of improving journalists' work lives more broadly:

But that's where our focus on the deployment of these tools in our newsrooms are, which is to create... except that this technology is real, and then figure out ways that it might improve our

work lives... reduce some drudgery, create a better output [where] the humans are still in charge.

In addition to news workers, the experts that were interviewed also agreed with the importance of having a human in the loop throughout the AI value chain. Jared Schroeder explains that there is less of an issue of AI eliminating any need for journalists to do their job, and more of an issue of AI necessitating a shift in business models. He illustrates this point using the AI news report tool Perplexity as an example:

Where will Perplexity or any AI tool get information if there's no journalists out there getting it? When I first saw Perplexity... I was like 'This is it. Journalism is over. Perplexity, this is it.' And one of my colleagues was like, 'Well, is Perplexity going out there and interviewing people? Is it out there checking the facts and holding people accountable?' And I was like 'no', and he's like, 'Well, then journalism's not done.' But someone's got to pay for it. That's the big problem there.

The extent to which AI in journalism will displace human labor is a point of contention among many experts and practitioners, with those interviewed for this study being no exception. Charlie Beckett presents the “optimistic case” for AI and labor as freeing up time for journalists:

...the optimistic case is it then frees up time and resources for [journalists] to do the human work, the stuff the AI is not so good at. And the business model case for that will be, because that is going to what makes your brand stand out.

He explains that one impact that has already been catalyzed by AI more generally, before the explosion of generative AI, is the demand for technical skills in the field:

...we've already seen with pre-generative AI that it was creating, obviously, a demand for tech savvy people: data scientists, computer scientists and so on programmers etc. And, of course, that's going to increase.

A shift brought on by the more recent, increased integration of generative AI into journalistic production is the creation of hybrid roles, which Beckett explains are new to the field:

...there's often been interesting hybrid careers because as I said before, it's always the human in the loop. You want to retain editorial oversight. Even with things like revenue. You want to be careful, you don't want the AI to go rogue if you're trying to get subscribers and so on. So, there are these interesting hybrid roles—Data Editor or Head of Product Development—these are terms, roles that weren't there five, ten years ago.

The creation of these roles presents two major challenges for newsrooms upon which respondents elaborated. The first challenge is the journalism industry needing to compete with the tech industry to attract talent to fill these roles, which Beckett describes as a “paradox” and a “real problem”:

So there's a bit of a paradox going on here: on the one hand, you're going to get people who are savvy about the tech. And there's a real problem here for the news industry to attract those

people when they'd rather much rather work for Google or Meta or OpenAI. But at the same time, you're actually going to need more of that human stuff.

The second challenge, as articulated by Ryan Restivo based on his experiences, is that of training people in newsrooms with different technical backgrounds to adapt to new roles and engage with new forms of software:

Through my experience, I have some very good stories about how you train people in a newsroom. And working with people in the newsroom... they do not all have similar technical backgrounds, they do not embrace new software as well as you may think. So it can be difficult to train new people to get them used to a new system, or something that is new.

Ultimately, in order to adapt to the new hybrid careers and roles brought on by the integration of AI into the field, journalists face the task to “embrace” new software, a term Restivo uses that evokes two components: a mindset and a skill set.

Overcoming these challenges can be rewarding, however. In the experience of Joe Amditis, he has seen AI provide a direct value add on the back end, allowing for journalists to free up time for what he sees as work more fundamental to the purpose of becoming a journalist in the first place:

It's cliché, but everybody talks about herding cats and wearing different hats and stuff like that. And it really puts a strain on a lot of these folks who just want to do journalism, they just want to

report and serve their community, and provide the critical information that their residents and community members need and want.

In Amditis' experience, the components of journalistic labor that entail “herding cats,” “wearing different hats” and otherwise assuming various roles, infringe upon the ability for news workers to “do journalism”. He therefore sees AI as a force of autonomy in terms of journalistic labor. Furthermore, he notes that the accessibility of many AI tools enables this autonomy at a low barrier to entry, freeing journalists from less meaningful, more time-consuming administrative tasks such as CMS, text reformatting, structuring email lists, etc.:

And so my biggest thing I see as the most one of the most valuable opportunities for these tools, especially since they are relatively and incredibly low cost... they're free to test and play around with—even the top models like GPT-4 or Bing in Creative Mode—that low opportunity, or educational costs, combined with the relatively low operational cost at the scale that they're going to be working on, is the biggest possibility.

At the same time, Amditis sees immense value in newsrooms that can harness the computational skill sets of workers to facilitate innovation:

...those kinds of skills, the ability to not just code or talk to computers if you do it like that, but the ability to have computers talk to each other and then talk to people who don't know computers about computers... be that middle ground that people can talk to; the people who can talk to computers.

In addition to the importance of hybrid roles themselves, Amditis emphasizes the crucial nature of communication between these roles as he espouses talking “to people who don’t know computers about computers”. On the topic of skill sets, David Cohn offers the analogy of a chainsaw to assert his understanding of the human element in journalistic labor:

I don't see generative AI as a panacea. In terms of labor, I do view it like the evolution from a handsaw to a chainsaw. So you ever see old videos, black and white videos of people in the 1800s? And two of them would use this big handsaw to cut down a tree. And then you can imagine one person with a chainsaw could do the same thing in a third of the time. So the important thing to point out is that it still needs an operator. You still need a person. The chainsaw doesn't operate itself.

Cohn notes the twofold parallel here between the chainsaw and generative AI: 1) a reduction in the amount of time needed to complete tasks and 2) the continued need for a human “operator”. He further expounds on the second point by explaining that the human operator must be uniquely skilled:

And the second thing to point out is: actually you need a more skilled operator because it's actually arguably more dangerous. A chainsaw can do a lot of damage immediately. So you need to be skilled, you still need someone and but that one person gets a lot more done.

He explains how, with the invention of the ATM, while larger banks lost employees, the net number of bank employees increased due to this reduced cost of opening new branches:

When the ATM was invented, all those bankers thought they were gonna lose their jobs. What ended up happening is that the cost of opening up new branches dropped... And as a result, while each one of those banks had less employees than the original one, there... were more bank employees than there were before.

Cohn compares this phenomenon with what he sees as the potential benefit of AI in terms of newsroom employees:

We have the New York Times and The Washington Post... we have these cathedrals, and they're great. I'm not bashing them. And maybe they need less people. But we have news deserts. And if the cost of opening up newsrooms drops, where before you needed 50 people, and you needed a big audience. And maybe you only need two people now. And so can two people open up in a news desert and actually be successful with generative AI?

While generative AI may reduce the labor force at larger news organizations such as the New York Times and the Washington Post, Cohn sees the impacts around him as being a potential boon for news coverage in news desert areas, presenting a potentially optimistic case for the future of AI and labor in journalism.

5.8 Value Alignment

Another common topic of discussion for respondents was values; in particular, the challenge of aligning the values of journalism with those of the tech industry. Practitioners in the technology and journalism spaces articulated the difference in value systems between the two industries, and how the integration of the two fields inherently necessitates an evaluation of the superimposition of these value systems upon each other.

Ryan Restivo explains the importance of interrogating tech company products in terms of how they are built, as they are not all built with the same ideas and goals in mind necessarily:

So even for all the crazy named companies, they're gonna have their own proprietary models that are basically just things that probably they lifted off of GPT models or something else. I think people...maybe they're not going to be as cynical as I am...but maybe they just need to be open to questioning, like how these companies that pitch people... how they built their stuff.

He explains that many technologies are overpromised in terms of their value and their sales pitches can obfuscate the extent to which their value directly helps people:

And be more open to pressing questions of... you have this cool thing that's gonna solve all my problems, but how does it actually work? And get a little bit more on that base understanding of how it works... My tool exists to help people and my tool is very open about how it's going to help people and I'm trying to be open about the results of what I'm trying to do to help people. I can't say that everyone else is going to do that.

Restivo explains that not every company or product owner will necessarily be “open” so deliberate interrogation of the technology itself is often warranted. The U.S.-based computer scientist and startup founder echoed this sentiment, explaining that there is a mix of honesty and ignorance in the AI startup and investment space:

You have the same thing from startup founders. Some inside investors are honest about it, you know, some of them are actually like, ‘Hey, this is gonna create some problems as well.’ So at this point, at this moment in time, it's willful ignorance to really claim that there are no negative consequences to this. Maybe two or three months ago, you could honestly believe that things are gonna go in the right direction, if you weren't deeply familiar with what's going on.

She goes on to explain that this “willful ignorance” is a product of tech companies’ business priorities, and the capitalistic drive for efficiency and profit:

And in business very often the priorities are make it cheaper, make it faster, save somebody's money. And they look at this as a way to accomplish that efficiency, saving money. That makes it the least bit about what they're doing to humanity, or to people who are gonna lose their jobs. It's really depressing.

The hype associated with AI around the time ChatGPT released coincided with a corresponding dip in cryptocurrency interest amidst a wave of a high profile scams and scandals; some respondents observed parallels of speculation and get-rich-quick schemes between industries, such as Joe Amditis:

I resisted for a long time, the sort of grifter, 'pay \$2 for this training, or the prompt book' stuff. It was very clear, that easy transition from crypto bro to AI grift. So I'm very hesitant... I know how quickly easy technology, whether it's drop-shipping or crypto or AI in this case, now appeals to that, sort of, 1000-yard entrepreneurial stare. So I tried to avoid that as much as possible.

The drive to innovate in Silicon Valley tech culture facilitates a hype machine that doesn't exist to the same extent in other industries. Damon Kiesow explains that industries such as tech and journalism diverge in their value systems due to differences in incentives:

...journalism has a certain set of logics, technology has a certain set of logics... doesn't mean one is better than the other. But they're not complimentary. They don't overlap effectively because tech is not concerned about ethical uses and communities and accuracy and public good in the way that journalism is.

He goes on to further articulate the lack of “overlap” in priorities between the two domains, contrasting the incentives of tech companies versus news outlets:

Tech is about entertainment. Distraction. Profit. News is about knowledge building and trust and reliability.

Kiesow explains that integrating a “set of logics” from one industry into another requires an evaluation of potential externalities resulting from such an integration:

So anytime you're adopting a technology that's basically a black box, and you're not interrogating the technology to make sure that the logics aren't imposed in a way that distorts the journalism logic, you're making a mistake. You may not know you made a mistake for two years [until] something happens that breaks the model so to speak, but you're imposing risk.

Interrogating the “black box” is an important point here, as it implies that the value systems embedded in these technologies are often hidden or opaque. On the journalistic side, however, it is still essential to obtain and maintain partnerships—and relationships more broadly—with platform companies and technologists in the interest of their business and audiences. As Jeff Jarvis points out, these relationships are essential even as journalists broaden their own technological skill sets:

If we want to advance our field of journalism, with the aid of technology, you need to work with the technologists still, even though we can now program in our language... even though there's a lot we can do.

Similarly, Kiesow contends that it is important to bring in technically oriented perspectives into newsrooms despite the divergence in value systems between the two industries. He cites the 2024 hiring by NPR of a new CEO from the world of technology – Katherine Maher, former head of the Wikimedia Foundation – as an example of this decision:

The key to this is having people who aren't just journalists running newsrooms or running the business. But... I was really interested in... the new CEO of NPR coming from Wikimedia. That's interesting. Is she gonna bring a mission-based but technologically mature point of view to the operation of that organization?

Ryan Restivo sees the conflict between news values and tech values directly. As the publisher of YESEO, he views the future of AI tool development as one with a forked road, with a path toward profit and a path toward helping users:

A lot of it is going to come down to how... a lot of companies that are going to sprout up—that are going to be in this space over time—are going to tell people how they help people get to their solutions and be more open and transparent about how they're using the tools that they're using to get there.

He believes that newsrooms should ultimately be cautious in discerning the aims of secretive companies:

Because I feel like the more secretive that companies are... they may be shooting for giant contracts and big funds. But that doesn't necessarily mean that they're in it for the right reasons, or that they're really their real aim is to help users, maybe this is to get a lot of money, and then kind of, you know, get out of here. But, I think that there are a lot of different things that, you know, newsrooms should be cautious about... and I guess my prediction, I guess, is that there'll be a lot more tools that pop up out of nowhere.

As “a lot more tools” emerge “out of nowhere,” newsrooms must ensure that they develop mutually beneficial relationships with platform companies offering these tools such that they harness the value of technological acumen while reconciling their own objectives with the incentives of these companies.

5.9 Summary

News workers and experts alike spoke to the shift that artificial intelligence has brought about to the realm of journalism, both at a systemic and individual level. Respondents discussed their experiences with “AI” as a hype machine, a tool, a value driver, a societal risk, an existential risk, an emancipatory innovation, and an inevitability. Across all themes, a central aspect was the existence of both challenges and benefits inherent to the integration of artificial intelligence systems into journalistic practice.

When considering *AI as a tool*, respondents found it helpful in an assistive capacity to journalists for a variety of tasks outside of the wholecloth generation of articles, such as transcription, translation, and data analysis. Respondents noted the useful nature of AI tools in pre-production, such as helping with verification, information gathering, and pre-writing tasks. However, the general consensus among respondents was that the “tool” functionality of AI was less practical in the actual writing of articles.

At the systemic level, respondents discussed how newsrooms both historically and today respond to technological *innovation*. Rapid advancements in AI, including generative AI, have generated both excitement and skepticism, resulting in mixed levels of adoption among newsrooms. Experts believe that AI has the potential to disrupt existing technologies and

transform news production; however, cautious approach to AI adoption, with concerns about hype, premature implementation, and being “the one that makes the next mistake” (Restivo).

Nonetheless, respondents expressed the importance of proactivity with AI, particularly in terms of newsroom *business models*. News workers spoke to the experience of shifting organizational business models through augmented content production and revenue generation. The value add of language models from the perspective of scale, speed, and accuracy in terms of its benefit to news organizations financially. However, some respondents also noted concerns about competition with search engines and the potential loss of search revenue. Such a concern underscores another important consideration that respondents emphasized: understanding their *audience* and leveraging this knowledge to optimize advertising and monetization strategies. Both news workers and experts noted that AI can assist in bridging the gap between newsrooms and audiences – whether they be local communities, subscribers, or users.

As audiences reckon with content that may or may not have been written using generative language models, respondents underscored the critical role that *trust and transparency* play in the AI integration process. Respondents agreed that news organizations must be transparent about the use of AI in their reporting to maintain credibility. Debacles of “stupidity” (Jarvis) will lead to public scrutiny, such as the backlash that Sports Illustrated received for failing to disclose their experimentation with generative AI in their reporting.

Newsrooms that choose to experiment with this new technology must not only be transparent, but also institute *guardrails* in order to mitigate risks. Both news workers and experts discussed the ethical implications of using AI in journalism, including bias in datasets, biased outputs, and the potential for misinformation campaigns. Despite rapid advancements in technological capabilities of AI, all of the respondents talked about the challenges of integrating

AI into newsrooms, and the corresponding importance of the *human element*. Humans are essential for making decisions, providing context, and ensuring ethical use of AI. Generally, respondents echoed many of the optimistic perspectives regarding journalistic concerns about the impact of AI on newsroom jobs, such as the fact that AI “frees up time” (Beckett) for journalists to pursue the more human aspects of their role.

Experts noted the importance of *value alignment* and collaboration between journalists and technologists in order to develop and use AI in a responsible and effective way, and news workers reflected on their experiences managing, working with, and otherwise interacting with practitioners with technical backgrounds. A common idea among respondents was the potential conflict between the values of journalism and the tech industry. These divergent values differ among organizations, and influences: a) the extent to which these organizations are willing to experiment with publishing AI-generated content and b) the level of human oversight applied during the process. In the next section, I discuss the results of a series of survey experiments designed to elicit reader perceptions on said content.

Chapter 6: Survey Analysis

6.1 Phase 1

After collecting the data for Phase 1 of the study, the responses to the “List 2-10 adjectives describing the article” question were filtered and cleaned (reduced to individual adjectives) to ensure usability and reliability. This process resulted in the extraction of 44 unique descriptive words (59 total) used by respondents to characterize the AI-generated news content. Subsequent refinement and consolidation yielded 26 unique descriptors for the final list. Of the 44 descriptors, some were synonymous (e.g. Amazing and Fabulous) and as such were redundant for use in the exploratory factor analysis. Furthermore, certain measures were eliminated from the study because they are not appropriate as descriptors of all news, only certain news content or specific sources of news. The list of the measures eliminated from the original 44 descriptors is shown below, along with the rationales for elimination:

1. Confrontational - appropriateness
(reaction to subject matter)
2. Decisive - redundancy (conclusive)
3. Definitive - redundancy (conclusive)
4. Digital - appropriateness (reaction to
subject matter - digital rights act
article)
5. Digital - appropriateness (digital
rights act article)
6. Disappointing - appropriateness
(reaction to subject matter)
7. Fabulous - redundancy (amazing)
8. Factual - duplicate
9. Flowing - changed to Coherent
10. Great - redundancy (amazing)
11. Greedy - appropriateness (reaction to
subject matter)
12. Informative - duplicate
13. Informative - duplicate

- | | |
|--|--|
| 14. Informative - duplicate | 24. Thought-Provoking - duplicate |
| 15. Informative - duplicate | 25. Threatening - appropriateness
(subject matter) |
| 16. Interesting - duplicate | 26. Unbelievable - redundancy
(believable) |
| 17. Lengthy - redundancy (long) | 27. Unbiased - duplicate |
| 18. Limiting - appropriateness (reaction
to subject matter) | 28. Uncompromising - appropriateness
(subject matter) |
| 19. Long - duplicate | 29. Unfair - redundancy |
| 20. Long - duplicate | 30. Uninteresting - redundancy |
| 21. New - redundancy (timely) | 31. Verbose - redundancy |
| 22. Possible - appropriateness (subj
matter) | 32. Wrong - redundancy (factual) |
| 23. Sad - appropriateness (subj matter) | |

Two other minor changes were made: first, the descriptor Flowing was changed to Coherent, a synonym, because the latter is a more widely understood term. Second, words with negative/opposite force were changed to their positive counterparts (Unbiased to Biased, Inconclusive to Conclusive) for the sake of clarity. The full list of descriptors used (bolded) along with the rationale for filtering given words (unbolded) is shown in the Appendix.

The final list of 26 words is shown below:

- | | |
|---------------|-------------|
| 1. Amazing | 4. Boring |
| 2. Believable | 5. Coherent |
| 3. Biased | 6. Complex |

- | | |
|-----------------|-----------------------|
| 7. Concise | 17. Moral |
| 8. Conclusive | 18. Neutral |
| 9. Detailed | 19. Precise |
| 10. Educational | 20. Rote |
| 11. Factual | 21. Technical |
| 12. Fair | 22. Thorough |
| 13. Honest | 23. Thought-provoking |
| 14. Informative | 24. Timely |
| 15. Interesting | 25. Weird |
| 16. Long | 26. Wordy |

These descriptors encompassed traditional news story attributes such as fairness and neutrality but also included novel descriptors like “Technical,” “Precise,” and “Weird.” This suggests that broader social perceptions of AI may influence article descriptions.

6.2 Phase 2

As previously mentioned, 261 responses were recorded for the Phase 2 survey. Each response included ratings of the 26 variables for three articles, for a total of 783 observations for each variable. Of the 20358 possible values in the data set, 35 were missing (NA). The analysis for Phase 2 was performed using R and RStudio, version 1.1.42, which are widely recognized for their robust statistical capabilities and flexibility in handling complex data sets. The correlation matrix was analyzed using the Pearson method, known for its efficacy in measuring the linear correlation between variables. For the estimation method, this study utilized the minimum

residual method, or minres. The method of estimating communalities, essential for understanding the shared variance in observed variables, involved using both h^2 and u^2 methods.

Communalities (h^2) represent the proportion of each variable's variance that can be explained by the factors, while uniqueness (u^2) represents the proportion of variance that is unique to the variable. Together, these methodological choices ensured a rigorous and comprehensive approach to uncovering the underlying dimensions of reader perceptions of AI-generated news content sourced in the Phase 2 survey.

Before proceeding with the Exploratory Factor Analysis (EFA), it was crucial to determine the suitability of the collected data for such statistical processing. Two key tests were employed for this purpose: Bartlett's Test of Sphericity and the Kaiser-Meyer-Olkin (KMO) Measure of Sampling Adequacy. Bartlett's Test of Sphericity is used to test the hypothesis that the variables are uncorrelated in the population. The test checks if the correlation matrix is an identity matrix, which would indicate that factor analysis is inappropriate. For the data, the Chi-square test statistic from Bartlett's test was 9189.899, with a p-value significantly less than 2.22×10^{-16} , essentially approaching zero. This extremely low p-value strongly rejects the null hypothesis of the test, suggesting that the variables are related and, therefore, suitable for factor analysis.

The KMO test was deployed to measure the adequacy of sampling and determine if the partial correlations among variables are small, which is preferable for factor analysis. The KMO statistic varies between 0 and 1, with a value closer to 1 indicating that factor analysis is appropriate. In this study, the overall KMO Measure of Sampling Adequacy was 0.92, far exceeding the minimum acceptable level of 0.50. A KMO value of 0.92 is considered excellent, indicating that a significant amount of variance might be explained by underlying factors and

that the data is very suitable for EFA. In sum, both tests strongly indicated that the data was appropriate for Exploratory Factor Analysis, ensuring the reliability and validity of the subsequent analysis.

6.2.1 Exploratory Factor Analysis

In the exploratory factor analysis of this study, two distinct factor rotation methods were employed: varimax and oblimin. Varimax, a widely used method, is an orthogonal rotation technique that simplifies the interpretation of factors by maximizing the variance of squared loadings of a factor on all the variables in a factor matrix. This method assumes that the factors are uncorrelated (orthogonal to each other) and thus focuses on clarifying which variables load highly on each factor. On the other hand, oblimin is an oblique rotation method that allows for correlations between factors. This is more realistic in many social science contexts where underlying constructs are often interrelated. The oblimin method can provide a more accurate representation of the data when factors are expected to be correlated.

The decision to include results from both varimax and oblimin rotations in the study was driven by the objective to gain a comprehensive understanding of the data. While orthogonal rotations like varimax simplify the structure and interpretation of factors, they might not always represent the true underlying relationships in the data, especially when factors are correlated. Oblique rotations like oblimin, although potentially more complex to interpret, can offer a more realistic picture of these inter-factor relationships. By examining the results of both rotation methods, the study leveraged the clarity and simplicity of orthogonal rotations and the realistic representation of factor correlations provided by oblique rotations. This dual approach allowed for a more nuanced and thorough exploration of the factor structure of reader perceptions of

AI-generated news, ensuring that the findings were both interpretable and closely aligned with the actual data structure.

In this study, the criterion for retaining factors was primarily based on eigenvalues, with additional consideration given to scree plot analyses (see Figs. 6.1 and 6.2 below). Factors with eigenvalues over 1 are typically considered significant as they explain more variance than a single observed variable. For the varimax rotation, the eigenvalues for the highest five factors were 4.67, 2.17, 1.89, 1.23, and 0.97. The fifth factor, with an eigenvalue slightly below 1, was also retained based on its proximity to 1 and the insights provided by the scree plot in Fig 6.1. Including the fifth factor was deemed important as it appeared to contribute meaningful information about the data structure relative to the next highest eigenvalue based on its positioning on the scree plot.

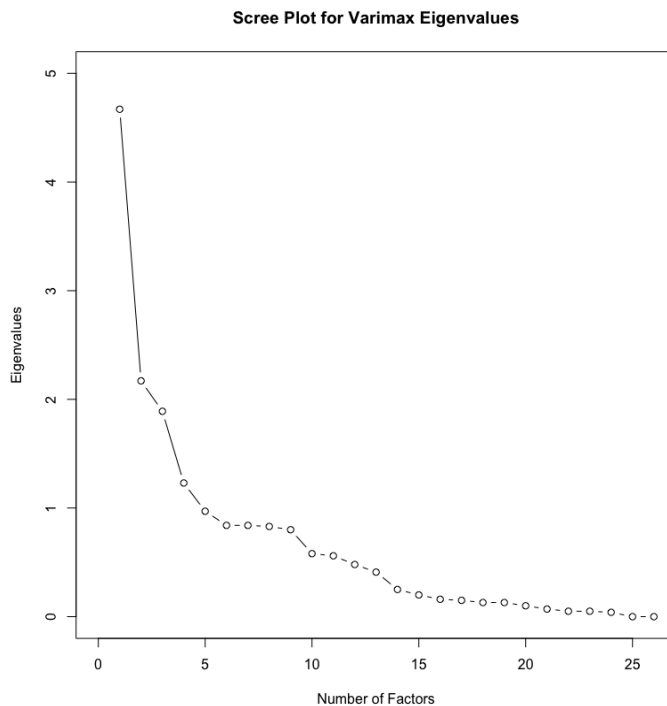


Figure 6.1: Scree Plot (Varimax Rotation)

In the oblimin rotation, which allows for correlation among factors, a total of nine factors were retained. The eigenvalues for these factors were 1.40, 1.38, 1.23, 1.09, 1.07, 1.01, 0.98, 0.95, and 0.94. Here, factors with eigenvalues close to but less than 1 were also included, as indicated by the scree plot analysis. This was based on the understanding that in oblique rotations, lower eigenvalues can still be meaningful due to the potential correlations between factors. Both eigenvalue criteria and scree plot (see Fig 6.2) analyses were instrumental in determining the number of factors to retain for each rotation method.

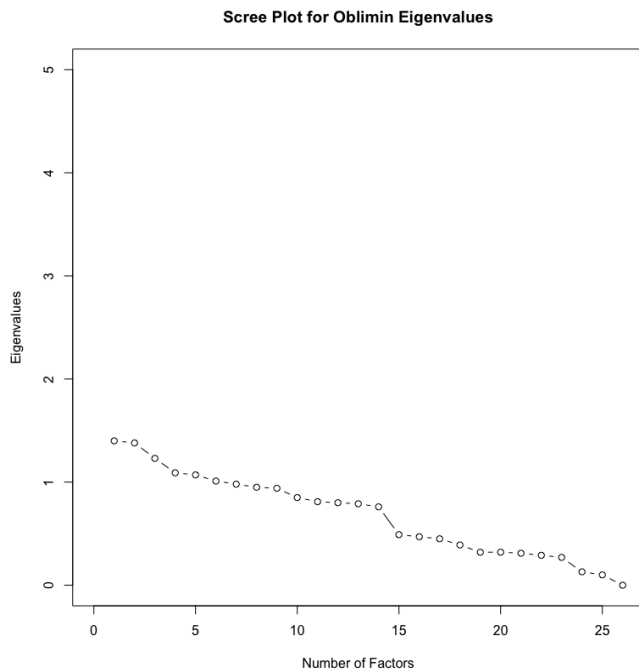


Figure 6.2: Scree Plot (Oblimin Rotation)

Tables 6.1 and 6.2 below summarize the eigenvalues for the factors retained under each analysis, including the proportion of variance and cumulative variance explained by each factor. Analysis of Table 6.1, the table for the varimax rotation with the given eigenvalues, indicates that the first factor explains a significantly larger portion of the variance in the data compared to

subsequent factors. The high eigenvalue in the first factor suggests that this factor in particular is the most influential in explaining the variability in the dataset. The first factor accounts for 27% of the variance, which is a substantial amount, indicating that this factor captures a significant portion of the information in the data set. The following factors, with proportions of 12%, 11%, 7%, and 6%, contribute progressively less to the total variance, but each still represents a meaningful aspect of the data. These five factors together explain 62% of the cumulative variance in the data. While this variability explained value is a substantial amount in social science research, it also illustrates that there is a moderate amount of variance in the data that is not captured by these factors, indicating other considerations might be influencing variance in the data.

Factor Number	Eigenvalues	Proportion of Variance	Cumulative Variance
1	4.67	0.27	0.27
2	2.17	0.12	0.39
3	1.89	0.11	0.50
4	1.23	0.07	0.57
5	0.97	0.06	0.62

Table 6.1: Eigenvalues and Variance Explained (Varimax Rotation)

As shown in Table 6.2 below, the eigenvalues obtained after the oblimin rotation range from 0.94 to 1.40 for the nine factors retained. The fact that these eigenvalues are relatively close to each other suggests that each factor contributes somewhat similarly to explaining the variance

in the data, unlike in the varimax rotation. Each factor contributes between 5% to 8% of the variance. This even distribution further reinforces that the data's underlying structure contains no single factor dominating the explanation of variance. The 0.58 cumulative variance explained by the nine factors is significant, but similar to the value in the varimax rotation, it also shows that 42% of the variance in the data is not explained by these factors.

Factor Number	Eigenvalues	Proportion of Variance	Cumulative Variance
1	1.40	0.08	0.08
2	1.38	0.08	0.16
3	1.23	0.07	0.23
4	1.09	0.06	0.29
5	1.07	0.06	0.35
6	1.01	0.06	0.42
7	0.98	0.06	0.48
8	0.95	0.05	0.53
9	0.94	0.05	0.58

Table 6.2: Eigenvalues and Variance Explained (Oblimin Rotation)

Below, in Tables 6.3 and 6.4, are the factor loadings in tabular format, along with the variables that load onto each salient factor, and labels for each salient factor. Salience was defined as a factor having loadings of above 0.30 from variables with their highest loading on said factor.

Table 6.3: Factor Loadings of News Stories Ratings - Varimax Rotation

Measure	Factor 1: Quality	Factor 2: Engagement	Factor 3: Alienation	Factor 4: Effort	Factor 5: Coherence*
Honest	0.83	0.16	-0.09	0.00	-0.05
Factual	0.74	0.16	-0.10	0.08	0.01
Fair	0.73	0.04	-0.15	-0.07	0.13
Believable	0.68	0.09	-0.21	0.00	0.15
Precise	0.68	0.34	-0.01	0.10	0.07
Thorough	0.65	0.30	-0.05	0.33	0.03
Informative	0.62	0.36	-0.19	0.25	0.04
Conclusive	0.56	0.31	0.04	0.13	-0.01
Concise	0.53	0.30	0.12	-0.19	0.21
Neutral	0.52	-0.08	0.06	-0.09	0.11

Detailed	0.50	0.34	-0.04	0.39	0.10
Moral	0.50	0.16	0.25	-0.07	0.00
Educational	0.49	0.39	0.04	0.22	0.05
Timely	0.45	0.28	-0.02	0.00	0.28
Interesting	0.32	0.74	0.06	-0.02	0.08
Thought-provoking	0.21	0.62	0.15	0.17	0.07
Amazing	0.17	0.54	0.35	0.20	-0.19
Weird	-0.23	0.06	0.73	0.10	-0.15
Rote	0.04	0.03	0.60	0.12	0.00
Biased	-0.13	0.21	0.58	0.24	0.02
Boring	-0.09	-0.52	0.52	0.32	0.09
Long	-0.05	0.00	0.38	0.69	-0.02
Wordy	-0.10	0.04	0.42	0.58	-0.15
Complex	0.15	0.24	0.28	0.39	0.18
Technical	0.28	0.25	0.30	0.34	-0.13

Coherent	0.45	0.01	-0.16	-0.04	0.56
----------	------	------	-------	-------	-------------

*Factor 5 had an eigenvalue of .97, which is under 1 but significantly higher than the eigenvalue of the next highest factor. Factor 5 also only loaded onto a single item (Coherent)

In Table 6.3, 18 of the 26 variables had a clearly high loading on one of the five factors (i.e. they possessed relatively small loadings on the remaining factors of at least .20 lesser than the highest loading in absolute value) whereas the other eight measures had moderate loadings on more than 1 factor.

6.2.2 Factor Labels - Varimax Rotation

- Factor 1 had high loadings of 14 variables, the highest of any of the factors by a wide margin. Variables such as Honest, Factual, and Fair, are directly tied to the credibility of the source and the message. Additionally, variables like Precise, Thorough, Informative, Conclusive, Concise, Detailed, Moral, Educational, and Timely point toward the overall quality and representativeness of content, encompassing both the depth and clarity of reporting as well as adherence to journalistic standards. As such, this factor was labeled “Quality,” reflecting aspects related to the quality of reporting/writing, perceived credibility of the source/message, and representativeness of journalistic norms.
- Factor 2 had high loadings of variables that evoke positive sentiment and a higher level of reader engagement. This factor included Interesting, Thought-provoking, and Amazing, all of which suggest content that is not only attention-grabbing but also stimulates deeper thinking and positive reactions from readers. As such, this factor was labeled “Engagement”.

- Factor 3 had high loadings of variables that suggest content might be off-putting, strange, or difficult to relate to. Variables like Weird, Rote, Biased, and Boring indicate content that either feels unfamiliar, one-sided, or lacking in engagement, potentially leading to a sense of alienation or disconnection for the reader. Thus, the factor was aptly labeled “Alienation”.
- Factor 4 had high loadings of variables such as Long, Wordy, Complex, and Technical, which point towards content that requires more effort to read and comprehend. These variables suggest that the articles might be verbose or complex, necessitating a higher level of effort from readers to parse through the information. This factor was labeled “Effort”.
- Factor 5 had a single high loading from the variable Coherent, indicating the logical flow and clarity of the content. This factor highlights the importance of coherence in how readers perceive and evaluate AI-generated news, with a focus on the logical structure and understandability of the content.

The analysis of secondary factor loadings, in addition to primary high loadings, offers deeper insights into the nuanced relationships between the higher level factors underlying the data. Among the variables that primarily loaded onto Quality, several also showed moderate loadings on Engagement. This suggests a relationship where content deemed high in quality (accurate, fair, believable) also engages readers. Specifically, attributes like being Precise, Conclusive, Concise, Thorough, Informative, Detailed, Educational, and Timely might not only contribute to the perceived quality of the content but also enhance its engaging nature. These

characteristics likely make the content more interesting and thought-provoking, drawing readers into a deeper engagement with the material.

Thorough, Informative, and Detailed variables having moderate loadings on Effort indicates that while these attributes enhance quality and engagement, they also require more cognitive effort from the readers. The depth and detail of the content might demand higher concentration and processing, thus impacting how effortless or taxing the reading experience is. The moderate loading of Moral on Alienation could be explained by the perception that discussions of morality in news content, especially when generated by AI, might create a sense of discomfort or ethical ambiguity. This could lead to feelings of alienation, as readers grapple with the AIs ability to handle complex moral and ethical issues.

In the Engagement category, Amazing having a moderate loading on Alienation suggests that while the content is engaging and perhaps surprising, it might also be perceived as overly sensational or unrealistic when attributed to AI, leading to a sense of alienation or skepticism. For variables under Alienation, the fact that Boring has an equally high negative loading on Engagement is quite telling. It underscores an inverse relationship where content perceived as boring not only fails to engage but actively disengages or alienates the audience. Additionally, its moderate loading on Effort implies that boring content might also be seen as requiring unnecessary or unfruitful effort to engage with.

All variables under Effort having moderate loadings on Alienation suggests a relationship where content that is long, wordy, complex, or technical could potentially alienate readers. This might be due to the increased effort required to understand such content, which could lead to frustration or disengagement, especially if readers do not immediately see the value or relevance of investing their time and cognitive resources. Finally, Coherent having a moderate loading on

Quality reinforces the idea that clarity and logical structure are not only crucial for understanding (coherence) but also contribute significantly to the perceived overall quality of the content. This underscores the importance of well-structured and logically consistent content in AI-generated news.

Table 6.4: Factor Loadings of News Stories Ratings - Oblimin Rotation

	Factor 1:	Factor 2:	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7*	Factor 8	Factor 9**
Honest	0.81	-0.01	0.05	0.03	0.03	0.02	0.00	0.06	0.08
Factual	0.71	0.00	-0.03	0.11	0.09	0.00	0.01	-0.02	-0.06
Believable	0.63	-0.09	-0.02	0.11	-0.01	0.00	0.05	-0.18	0.00
Fair	0.54	-0.04	0.08	-0.08	0.01	-0.10	0.10	-0.08	0.32
Moral	0.42	0.06	0.05	-0.04	0.14	0.28	-0.05	0.05	0.15
Long	-0.04	0.86	-0.05	-0.05	-0.01	-0.04	0.07	-0.05	0.02
Wordy	0.00	0.70	-0.05	0.06	-0.10	0.11	-0.05	0.11	-0.02

Boring	-0.02	0.17	-0.73	0.04	-0.01	0.20	0.10	-0.01	0.04
Interesting	0.04	0.01	0.61	0.16	0.08	0.22	0.16	-0.01	0.06
Thought-provoking	-0.04	0.14	0.39	0.26	0.10	0.21	0.12	-0.02	-0.04
Amazing	0.10	0.24	0.33	0.08	0.14	0.24	0.05	0.26	-0.08
Educational	-0.03	-0.04	0.01	0.71	0.03	0.05	0.06	0.03	0.09
Informative	0.24	0.01	0.05	0.62	0.00	-0.10	0.01	-0.05	-0.02
Concise	0.02	-0.17	0.04	-0.02	0.71	0.11	0.01	-0.08	0.07
Precise	0.22	0.07	0.08	0.15	0.48	-0.11	0.00	0.00	0.04
Conclusive	0.14	0.10	0.03	0.25	0.41	-0.06	-0.09	0.07	0.02
Thorough	0.23	0.22	0.06	0.20	0.32	-0.26	0.09	0.04	0.05
Detailed	0.07	0.26	0.11	0.20	0.30	-0.26	0.17	-0.03	0.02
Biased	0.05	0.22	-0.05	-0.02	0.06	0.53	0.15	0.02	-0.23
Weird	-0.22	0.17	-0.06	-0.13	0.07	0.39	0.11	0.33	0.16
Rote	-0.06	0.01	-0.28	0.13	0.15	0.38	0.11	0.17	0.04

Complex	0.00	0.02	0.01	0.00	-0.06	0.01	0.81	-0.03	-0.02
Coherent	0.12	-0.01	-0.06	0.03	0.25	0.00	0.09	-0.56	0.09
Technical	0.10	0.00	-0.06	0.11	0.14	-0.10	0.41	0.41	0.00
Neutral	0.08	0.03	-0.05	0.07	0.02	-0.03	-0.04	0.00	0.69
Timely	0.00	0.02	0.18	0.23	0.11	0.07	0.08	-0.25	0.26

*Factor 7 loaded onto one item (Complexity)

**Factor 9 loaded onto one item (Neutrality)

After applying an oblimin rather than a varimax rotation, resulting in the data in Table 6.4, 16 of the 26 variables had a clearly high loading on one of the nine factors (a difference of at least 0.20 between the highest and second-highest absolute values of each loading from the variable).

Conversely, 9 of the 26 factors had a moderately high loading in addition to a clearly high loading. One variable (Timely) failed to saliently load onto any factor.

6.2.3 Factor Labels - Oblimin Rotation

- Factor 1 has high loadings of five variables integral to the trustworthiness and reliability of news content: Honest, Factual, Believable, Fair, and Moral. This factor reflects attributes key to establishing trust in news. These variables suggest that readers place high value on authenticity, accuracy, and ethical considerations in news content, especially when generated by AI. This factor is aptly labeled “Trust”.

- Factor 2 has high loadings of variables Long and Wordy. This factor captures the verbosity and lengthiness of content. It reflects reader perceptions of articles that may be overly detailed or extended, potentially affecting the readability and accessibility of the news. This factor is thus named “Prolixity”.
- Factor 3 had high loadings of four variables representing the ability of the content to engage or disengage the reader. This factor highlights the importance of keeping the audience captivated and mentally stimulated, leading to its label as “Engagement” similar to the factor in the varimax results.
- Factor 4, having high loadings of the two variables Educational and Informative, emphasizes the informative value and educational quality of the news content. It reflects a preference for content that is enriching and enlightening, and is therefore labeled “Substance.” The term Substance encapsulates the depth, richness, and informative nature of the content, aligning well with the two variables that load onto the factor.
- Factor 5 possessed high loadings of the variables Concise, Precise, Conclusive, Thorough, and Detailed. This fifth factor pertains to the craftsmanship of the writing, focusing on clarity, accuracy, and depth in news presentation, and leading to its designation as “Clarity.” This label emphasizes the clear, concise, and precise nature of the writing style.
- Factor 6, having high loadings of the variables Biased, Weird, and Rote, reflects elements that might create a sense of estrangement or disconnection for the reader. It is appropriately termed “Alienation,” and is similar in structure to the factor from the varimax analysis.

- Factor 7, having a high loading solely of the variable Complex, captures the intricacy or sophistication of the content. This factor indicates the influence of complexity on understanding and engagement with the news and is labeled “Complexity.”
- Factor 8 had its highest loadings on Coherent and Technical. This factor relates to the mental effort required by readers to comprehend and engage with the content. Emphasizing clarity and technicality in news presentation, this factor is named “Effort,” and is constructed similarly to the factor in the varimax analysis.
- Factor 9, consisting of just Neutral, points to the impartiality and unbiased nature of the news. The factor is thus called “Neutrality.”

As previously mentioned, the analysis of secondary factor loadings, in addition to primary high loadings, offers deeper insights into the nuanced relationships between the higher level factors underlying the data. Among the variables with high Trust loadings, the additional moderate loading of Moral onto Alienation is notable in its consistency with the varimax results. This loading again suggests a complex relationship between ethical considerations and feelings of disconnection. When AI-generated content addresses moral issues, it might raise concerns or skepticism among readers about the AIs ability to navigate complex ethical landscapes, potentially leading to alienation. For the Engagement loadings, Thought-provoking loading moderately onto Substance and Alienation indicates that while engaging content stimulates deeper thinking, it might also touch on complex or sensitive topics that can alienate some readers. Amazing showing moderate loadings on Prolivity, Alienation, and Effort is intriguing. It suggests that while such content is captivating, it might also be perceived as verbose, potentially

alienating or requiring more cognitive effort to process. The similar loading onto Alienation in the varimax analysis reinforces this idea.

For the Clarity loadings, Conclusive, Thorough, and Detailed having moderate loadings on Substance align with the notion that clear writing often goes hand-in-hand with substantive content. The moderate loading of Thorough and Detailed on Prolixity and their negative loading on Alienation suggest a balance between depth and accessibility; while thorough and detailed reporting is valued for its substance, there is a risk of it becoming verbose and potentially alienating if not presented clearly. For the Alienation loadings, Weird loading moderately on Effort and negatively on Trust might reflect a perception that unconventional or unusual content, while intriguing, can be challenging to comprehend and might undermine the perceived credibility of the content. Rote showing a negative loading on Engagement suggests that content perceived as mundane or formulaic is not only alienating but also fails to engage readers effectively.

For the Effort loadings, Technical showing an equal loading on Complexity indicates a close relationship between the technical nature of content and its complexity, both contributing to the effort required in understanding the material. Timely loading moderately on Substance but negatively on Effort could suggest that while timely content is valued for its relevance and substance, it might be presented in a way that requires less cognitive effort, perhaps due to its immediacy or the nature of its presentation.

An oblimin rotation assumes factors are not independent and are correlated; as such, the strength of correlations between factors are a worthy subject of analysis. Table 6.5 shows the values of the factor correlations.

Table 6.5: Factor Correlations - Oblimin Rotation

	Factor 1: Trust	Factor 2: Prolivity	Factor 3: Engagement	Factor 4: Substance	Factor 5: Clarity	Factor 6: Alienation	Factor 7: Complexity	Factor 8: Effort	Factor 9: Neutrality
Factor 1: Trust	1.00	-0.04	0.30	0.62	0.57	-0.20	0.21	-0.23	0.42
Factor 2: Prolivity	-0.04	1.00	-0.17	0.24	0.11	0.29	0.48	0.34	-0.07
Factor 3: Engagement	0.30	-0.17	1.00	0.38	0.31	-0.08	0.06	-0.05	0.04
Factor 4: Substance	0.62	0.24	0.38	1.00	0.56	-0.03	0.42	-0.01	0.19
Factor 5: Clarity	0.57	0.11	0.31	0.56	1.00	0.11	0.35	-0.03	0.35
Factor 6: Alienation	-0.20	0.29	-0.08	-0.03	0.11	1.00	0.24	0.29	-0.04
Factor 7: Complexity	0.21	0.48	0.06	0.42	0.35	0.24	1.00	0.12	0.10
Factor 8: Effort	-0.23	0.34	-0.05	-0.01	-0.03	0.29	0.12	1.00	-0.14
Factor 9: Neutrality	0.42	-0.07	0.04	0.19	0.35	-0.04	0.10	-0.14	1.00

The two strongest correlations were Factor 1 (Trust)-Factor 4 (Substance) and Factor 1 (Trust)-Factor 5 (Clarity). This is logically consistent with the varimax rotation, given the combination of Trust-themed adjectives and quality-themed adjectives were combined in the first factor from that analysis. The next-strongest correlation was between Factor 4 (Substance) and Factor 5 (Clarity) themselves.

The Alienation factor emerged as a particularly salient and inter-correlative element in the analysis, underscoring its significance in shaping reader perceptions of AI-generated news

content. A notable trend observed was the common occurrence of variables with high loadings across multiple factors also sharing moderate loadings on the Alienation factor. This pattern suggests that while certain attributes of the news content are positively associated with factors like Trust, Engagement, or Clarity, they simultaneously hold the potential to alienate readers, perhaps due to the unconventional or off-putting nature of the content. The moderate negative loading of certain variables onto Trust in the obliquely rotated analysis further accentuates this point. It indicates the presence of a negative correlation between content generated by AI and the way in which its content contributes to feelings of alienation. These feelings could be potentially due to perceived biases or the impersonal nature in which AI-generated content reads (or primes to read) to consumers.

The finding that the Biased variable loaded highest onto the Alienation factor—rather than a factor typically associated with credibility or quality—in both factor structures (varimax and oblimin) is intriguing. There are two different considerations that might explain this phenomenon; the first being the influence of social AI perceptions on content perceptions, and the second being specific source- and content-based cues that relate to alienation. When readers know that a news article is generated by AI, their perception of bias could be influenced by preconceived notions about artificial intelligence. The loading of Biased onto Alienation might reflect broader societal concerns about the role of AI in media and information dissemination. As AI becomes more prevalent in journalism, there are growing discussions and apprehensions about AI's role in shaping narratives or perpetuating biases. This societal context could influence how readers interpret and react to AI-generated content, particularly regarding bias. There is often a skepticism about AI's ability to be truly neutral or unbiased, as AI systems can inadvertently reflect the biases present in their training data. This skepticism could lead to a

heightened sensitivity to any perceived bias in AI-generated articles, which might contribute to a feeling of alienation rather than just being a mark against perceived quality.

The concept of alienation in this context could be tied not just to the content of the news but also to its source. If readers perceive AI-generated content as inherently biased, regardless of its actual neutrality or balance, this perception could lead to a sense of disconnect or mistrust. The “machine” behind the news might be seen as less capable of fair and balanced reporting compared to a human journalist, contributing to a sense of alienation. The label of bias in a news article, especially one attributed to AI, might trigger stronger emotional reactions compared to other quality-related concerns. This reaction could lead to feelings of discomfort, distrust, or disagreement, aligning more closely with the theme of alienation. In this sense, bias isn't just a marker of quality but becomes a barrier to reader engagement with the content.

6.3 Phase 3

Based on the results from Phase 2 and previous survey research on reader perceptions of automated content, the following hypotheses were formed to be tested in the Phase 3 analysis:

H1a: Readers rate human-written news higher than AI-written news on engagement, substance, and clarity

H1b: Readers rate AI-written news higher than human-written news on prolixity, alienation, complexity, and effort

H1c: No difference will be observed between AI and human authorship on trust and neutrality

H2: Reader perceptions are consistent across topics

H1a was proposed based on previous studies in which survey respondents ranked human-written news articles higher than AI-written articles on similar metrics. Similarly, H1b

was proposed based on previous studies in which survey respondents ranked AI-written news articles higher than human-written news articles on similar metrics. H1c was proposed based on previous studies that found no significant difference between human- and AI-generated articles on similar metrics. H2 was proposed based on results from previous studies; the majority of studies that varied article topic as an independent variable found few to no significant differences across topics.

Data for Phase 3 was collected from respondents to compare how news readers in the United States perceive AI (ChatGPT)-generated and human-written journalism in regard to the underlying factors obtained in the oblimin-rotated analysis from Phase 2. Data was cleaned and refined such that the final set of values contained 153 responses in which respondents gave nine ratings (corresponding to the nine factors obtained in the oblimin-rotated analysis) from 1 to 5 or each of the three articles in response to the following question: “In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)”. A series of two-way, repeated measures ANOVA analyses were used to test the hypotheses above. As previously mentioned, the oblimin rotation provides a more accurate representation of the underlying data structure, particularly given the nature of the constructs being studied. Using the factors derived from the oblimin rotation in ANOVA analyses ensures that the analysis considers the potential interactions and correlations between factors, which could be crucial for addressing *RQ3: “How do news readers in the United States perceive automated news content and human-written news content relative to the descriptors derived from RQ2?”* effectively. Furthermore, given the assumption that the oblimin-rotated analysis more accurately captures the underlying structure of the data, using it will lead to more reliable and valid results in subsequent analyses, even beyond the ANOVA results summarized below.

H1a predicted that, in terms of engagement, substance, and clarity, reader ratings of human-written articles would be higher than the AI-written article ratings. In terms of engagement, reader ratings of human-written articles were actually significantly lower than those of AI-written articles ($p = 0.003$, $F = 8.67$); however, ratings of human-written articles were only lower for the finance and tech article ratings (see Table 6.6). A significant interaction effect was observed between article authorship and article topic on the engagement outcome variable ($p = 0.008$, $F = 4.89$). This suggests that for certain topics (such as finance and tech), respondents may perceive and engage with AI-authored and human-authored articles differently, while for other topics (such as politics), such differences in perception may not exist. The same was true for clarity: ratings of human-written articles were significantly lower than ratings of AI-written articles ($p = 0.004$, $F = 8.34$), but ratings were only lower for the finance and tech human-written articles (see Table 6.7). Again, a significant interaction effect was observed between article authorship and article topic on the clarity outcome variable ($p = 0.011$, $F = 4.58$). No significant difference was found between human- and AI-written articles for ratings of substance ($p = 0.90$, $F = 0.02$) (see Table 6.8). Therefore, H1a was not supported by the data.

Table 6.6: Engagement Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	3.70, 0.98	3.42, 1.00	3.37, 1.00
Human	3.13, 0.91	3.52, 0.93	3.15, 0.98

Table 6.7: Substance Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
-------------------	----------------	-----------------	-------------

AI	3.69, 0.95	3.58, 0.97	3.55, 0.88
Human	3.82, 0.80	3.61, 0.89	3.38, 0.81

Table 6.8: Clarity Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	4.08, 0.75	3.73, 1.07	3.80, 0.90
Human	3.56, 0.83	3.82, 0.88	3.59, 0.99

H1b predicted that, in terms of prolixity, alienation, complexity, and effort, reader ratings of human-written articles would be lower than their ratings of AI-written articles. In terms of prolixity, human-written articles were indeed rated significantly lower ($p < 0.001$, $F = 12.01$) than AI-written articles across all topics (see Table 6.9). In terms of alienation, no significant difference was found ($p = 0.57$, $F = 0.32$) between the two authorship groups (see Table 6.10). In terms of complexity, no significant main effect of authorship was found ($p = 0.55$, $F = 0.36$), suggesting that, on average, respondents do not differentiate between AI-authored and human-authored articles in terms of how complex they find them (see Table 6.11). This indicates that the overall impact of the authorship type—whether an article is written by an AI or a human—is not sufficient on its own to affect complexity perceptions across all topics. However, a significant interaction effect ($p = .008$, $F = 4.86$) was found between article authorship and article topic on the complexity outcome variable, indicating that the impact of authorship on complexity perceptions is contingent on the specific topic of the article. The same was true for effort: no significant main effect of authorship was found (see Table 6.12). However, a

significant interaction effect was found ($p = 0.02$, $F = 3.95$) between authorship and topic on effort. Therefore, H1b was partially supported by the data.

Table 6.9 Prolixity Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	3.08, 1.05	3.20, 1.24	3.06, 1.26
Human	3.04, 1.07	2.62, 1.15	2.72, 1.09

Table 6.10 Alienation Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	2.15, 1.20	2.31, 1.19	2.32, 1.14
Human	2.29, 1.14	2.29, 1.14	2.24, 1.17

Table 6.11 Complexity Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	3.10, 0.97	3.21, 0.98	3.11, 1.02
Human	3.43, 0.94	2.85, 1.12	3.04, 1.05

Table 6.12 Effort Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	3.38, 1.01	3.25, 1.17	3.39, 1.14
Human	3.70, 0.98	2.96, 1.04	3.21, 1.10

H1c predicted that, in terms of trust and neutrality, no significant differences would be observed between authorship types. In terms of trust, while AI-written articles were rated slightly higher on average across all article topics, the difference between the two authorship groups was not significant ($p = 0.11$, $F = 2.53$) (see Table 6.13). In terms of neutrality, a significant difference was found between authorship groups: on average, readers perceive AI-written articles to be more neutral than human-written articles ($p = 0.044$, $F = 4.07$) (see Table 6.14). Therefore, H1c was partially supported by the data.

Table 6.13 Trust Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	3.96, 0.78	3.90, 0.90	3.92, 0.81
Human	3.82, 0.74	3.84, 0.87	3.74, 0.93

Table 6.14 Neutrality Ratings (Mean, SD) by Authorship and Topic

Authorship	Finance	Politics	Tech
AI	3.76, 0.93	3.72, 1.11	3.73, 0.96
Human	3.34, 1.03	3.59, 0.85	3.77, 0.96

H2 predicted that reader perceptions would be consistent across topics. This hypothesis implies that the mean ratings for each outcome variable should not significantly differ among topics, regardless of the authorship. This can be evaluated by examining both the main effect of the topic and the interaction between topic and authorship. When controlling for the repeated-measures nature of the study, no significant main effects of the article topic were found

on any of the outcome variables. However, as previously mentioned, statistically significant interaction effects were found between authorship and topic for four of the nine outcome variables: complexity ($p = 0.008$, $F = 4.86$), effort ($p = 0.02$, $F = 3.95$), clarity ($p = 0.01$, $F = 4.58$), and engagement ($p = 0.008$, $F = 4.89$). While overall perceptions might not differ across topics, the way in which they are influenced by authorship does differ. This finding does not fully support the hypothesis that perceptions are consistent across topics because it indicates that perceptions are dependent on the combination of topic and authorship. Therefore, H2 was only partially supported by the data.

The figures below display model estimates for the effects of authorship on the nine outcome variables within each of the three topic areas: finance (see Figure 6.3), politics (see Figure 6.4), and tech (see Figure 6.5). These visualizations allow for a deeper understanding of the interaction effects under which the effect of authorship varies by topic by making these patterns more discernible. For the finance articles, human authorship was rated significantly lower than AI authorship in terms of engagement, clarity, and neutrality, but was rated significantly higher in terms of complexity. For the politics articles, human authorship was rated significantly lower in terms of prolixity and complexity. For the tech articles, no significant differences were observed. The coefficient values for each of the outcome variables in the finance articles were most variable, whereas the coefficient values for the tech articles were closest together. The coefficient values for the tech articles were also closest to zero, suggesting that respondents perceived the AI-written and human-written tech articles to be more similar to each other than the articles with other topics. Across topics, variability in coefficient directions for most outcome variables were observed: the coefficient directions for engagement, substance, clarity, alienation, complexity, effort, and neutrality all differed across topics. Altogether, these

results suggest that perceptions of outcome metrics were partially contingent on the topic of the article respondents read.

Figure 6.3 Model estimates (finance topic only)

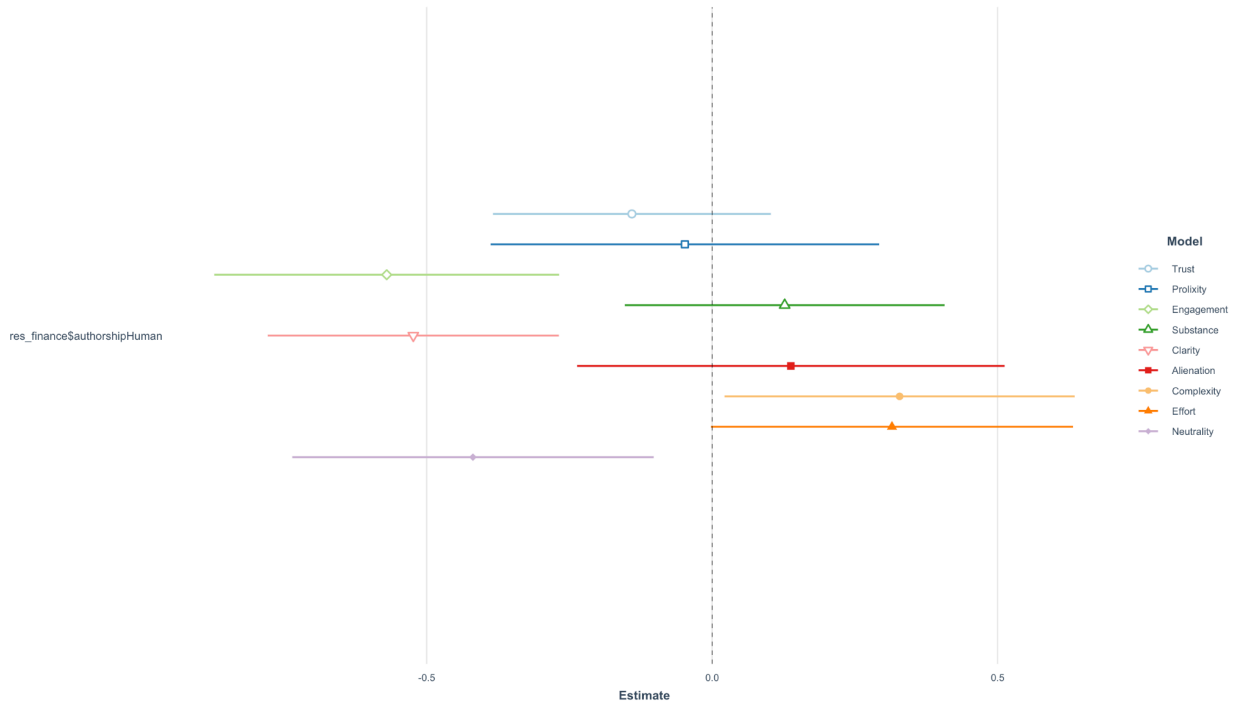


Figure 6.4 Model estimates (politics topic only)

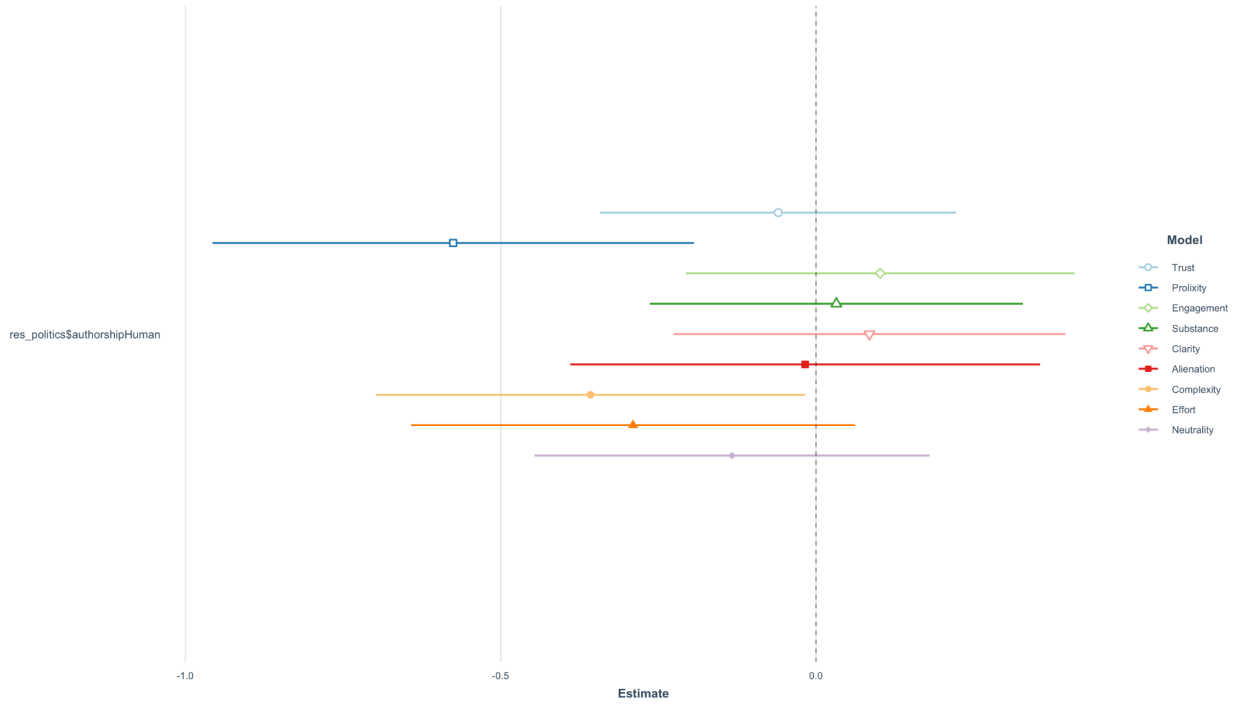
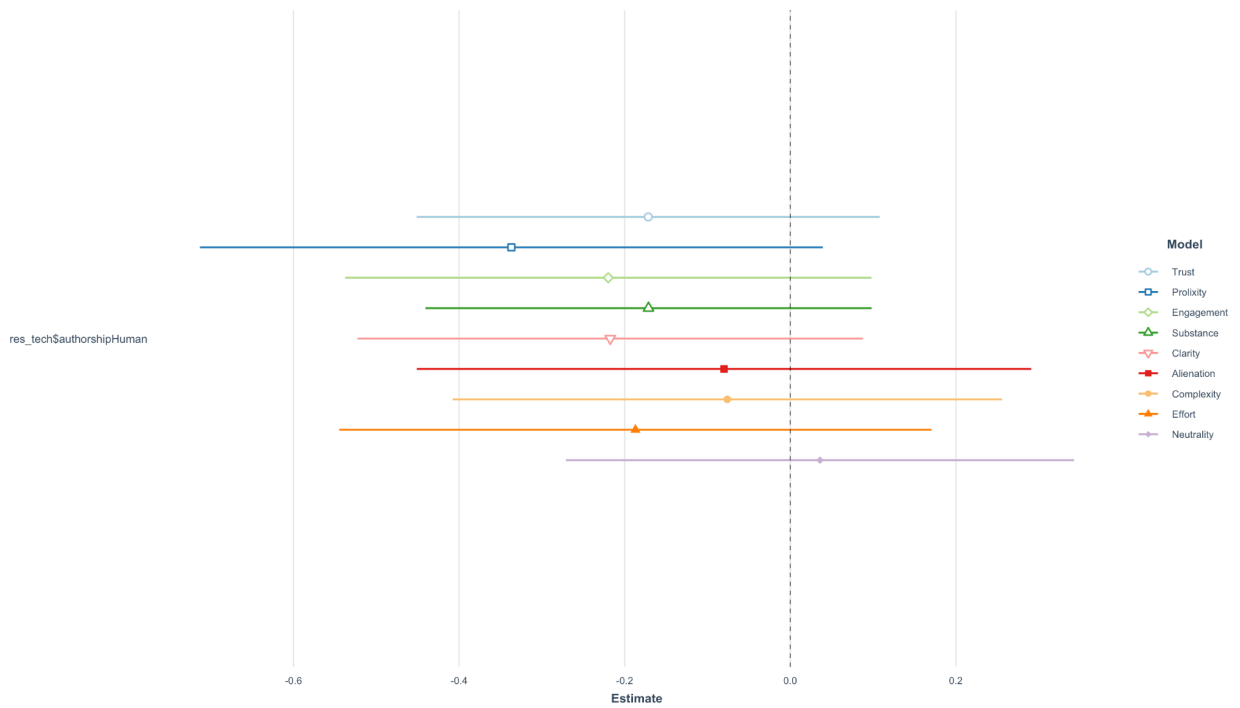


Figure 6.5 Model estimates (tech topic only)



Chapter 7: Discussion and Conclusion

This dissertation has explored the novel phenomenon of automated text generation in journalism from both a production and consumption perspective. Building on existing academic research that has delved into both perspectives, this analysis augments previous studies by addressing the theoretical challenges that an ontology of algorithms presents to the system of representation in which traditional journalism is deeply entrenched. Where human-written journalism can represent opinions and ideas that belie reasoning behind word choice and sentence construction, algorithmically generated texts arrange and render outputs that lack these explanations. Furthermore, AI-generated text in news articles challenges the expectations of readers who, by default, approach written work with the anticipation of interpreting representations that reflect intentional human thought rather than probabilistic inferences.

The interviews sourced diverse perspectives of a variety of individuals involved in the production of automated news content. Assessing the experiences of both news workers and experts garnered insights on how those who operate at the intersection of technology and journalism are reckoning with technologies like ChatGPT and Perplexity. Furthermore, these conversations offered an understanding of the systemic and individual-level changes catalyzed by the innovations in addition to the implications of these changes on a variety of stakeholders. The surveys complemented this assessment by evaluating the consumption aspect of journalism, revealing corresponding themes and trends among news readers presented with AI-generated news content. In addition to measuring the psychological underpinnings that govern media consumption in the first two phases, the survey component of the dissertation appraised reader distinctions between AI-generated and human-authored news through a comparative assessment in the third phase. In this final analysis, I address some of the considerations raised throughout

each component of the study, some limitations of the study overall, and theoretical implications of the study given the future of artificial intelligence in news production and consumption.

7.1 Considerations

7.1.1 Technology, Analysis, Mythology

Big Data as a cultural and technological phenomenon informed all aspects of the study. At a high level, this dissertation has explored how the age of Big Data, characterized by vast and complex datasets, has significantly impacted the field of journalism, leading to the increasing use of automation and AI in news production. This phenomenon forms the basis for the study's purpose, which is to examine the implications of this shift, focusing on both the production and consumption of AI-generated news. Big Data, driven by technological advancements, the ability to analyze vast datasets for patterns, and the perception of data as a source of superior knowledge, has led to a deluge of information. Consequently, the availability of data to train large language models has upended traditional news production methods, creating a demand for more efficient and scalable solutions and ultimately driving the integration of automation and AI (via these LLMs) into journalism. Responses from both interviews and surveys raised considerations regarding the interplay of these forces in automated news production and consumption.

The interviews emphasized the role of technology in gathering, analyzing, and processing large datasets, mirroring the first aspect of Big Data described by boyd and Crawford. Respondents consistently highlighted the practical applications of AI as a tool for journalists, particularly in tasks that involve managing and analyzing vast amounts of information. These

tasks include transcription, translation, data analysis, verification, and identifying patterns in data. Beyond the application of generative AI as a user-facing text generation engine, the computational power of large language models drives change and innovation across the journalistic value chain, as evidenced throughout the interviews. Respondents consistently highlight the practical applications of AI as a tool for journalists, particularly in tasks that involve managing and analyzing vast amounts of information. For example, the "unsexy AI," as Ryan Restivo calls it, used in back-end processes such as keyword extraction and data structuring demonstrates the integration of technology into existing journalistic workflows to handle the influx of information in the Big Data age.

The second aspect of Big Data mentioned by boyd and Crawford, analysis—drawing on large datasets to identify patterns and make claims—is evident in how interview respondents perceive the potential of AI to transform news production and business models. Respondents discussed how AI can help produce what Chris Dinn describes as "better content and more content" by automating tasks, identifying trends, and enabling the creation of personalized news experiences. Through the use of data analysis to identify patterns and cater to specific audience segments, generative AI offers news organizations a unique value proposition that was particularly salient in responses from interview respondents on the business side of news production.

Researchers and practitioners acknowledged the importance of data-driven insights in shaping content strategies, optimizing news production, and understanding audience preferences, demonstrating the practical application of data analysis in the field. This sentiment was not only evident in the survey component, in which audience preferences were directly sourced, but it was also particularly relevant among the portion of respondents who described audiences as

addressable economic markets. Ryan Restivo's invocation of a "feedback loop" reflects the initiative of those on the production side of news to continuously adjust to the evolving preferences of news consumers based on analyzing the data that consumers provide. As Jay Allred put it:

...we invite every reader to give us feedback on the sports articles that we publish, like every article that publishes has a link to a Google form where our readers can give us feedback.

Interview responses reinforce the mythological aspect of Big Data as boyd and Crawford explain it, particularly the belief in AI's potential to surpass human capabilities and provide a higher form of knowledge. The hype surrounding AI, especially after the release of ChatGPT, reflects the widespread fascination with AI's potential to revolutionize various fields, including journalism. Whether labeled "a very exciting, dynamic topic" or "overhyped a little," ChatGPT as an archetypal generative AI tool has clearly made an impact in terms of influencing social perceptions. In Jared Schroeder's experience with LLMs:

We're training indiscriminately, really lightly discriminately training on millions of billions like, we cannot conceive the amount of information that they're training these tools for. It's like, the entire history of humankind's knowledge is being pushed into these things. So the technical argument of the... I'm still standing on a very theoretical, like, not technical level, but the technical ability of these tools is like, there's never been anything like it.

In the survey component, the phases of the survey in which respondents were informed of AI authorship (Phases 1 and 2) assessed the latent impacts of these social perceptions. The 26 Phase 1 descriptors, sourced as adjectives describing the AI-generated article in the survey, encompassed traditional news story attributes such as fairness and neutrality but also included novel descriptors like “Technical,” “Precise,” and “Weird.”. These descriptors align with interview respondent perspectives on the hype of user-facing large language models – technical in their display of computational power, precise in their parametric reproduction of their corpuses, and weird in their uncanny ability to imitate texts written by humans. Ultimately, these new descriptors obtained in Phase 1 of the survey component mirror the three interconnected dimensions of Big Data–technology, analysis, and mythology.

The Phase 2 and 3 survey results further support this interplay. In the varimax-rotated exploratory analysis of Phase 2 data, the emergence of factors like Quality, Engagement, and Alienation reflects the influence of both technological capabilities and social perceptions in how readers respond to AI-generated news. The Quality factor, encompassing attributes like accuracy, fairness, and credibility, aligns with the desire for reliable information in the Big Data age, while Engagement and Alienation highlight the importance of how AI-generated content connects (or fails to connect) with readers on an emotional and intellectual level. In the Phase 3 analysis, the finding that AI-written articles were perceived as more neutral than human-written articles potentially reflects the mythology surrounding AI objectivity. Such a finding aligns with the machine heuristic, the rule of thumb that machines are more secure and objective than humans (Sundar and Kim, 2019). However, the lack of significant differences in trust ratings suggests that readers may not blindly trust AI-generated content and still value human judgment.

7.1.2 Trust and Transparency

The interview and survey results presented in this study reveal several key challenges concerning the integration of AI into newsrooms, particularly regarding trust and transparency. These challenges underscore the complexities of navigating this new technological landscape while upholding journalistic values and maintaining public trust. As discussed in Chapter 3, the ontological world of the algorithm changes the expectation of readership and readerly reception that are available to readers due to a shift in the locus of trust brought on by these evolving technologies. The results of this dissertation call attention to a transformation in the capacity, necessity, and locus of trust assignment. As a product of the human mind, textual language is linked to the consciousness and attentive effort of the writer. But as a product of algorithms, in which outputs are based on probabilistic arrangements and dispositions rather than direct representations, such traditional assignments of trust become untenable.

Trust arose as a salient theme in the interview analysis in terms of its assignment to journalists, news outlets, and AI as a socially constructed idea. For practitioners, trust and transparency often went hand in hand: being transparent about their news production process, authorship, and data management facilitated trust among readers, users, and/or customers. Similarly, practitioners found that trust was an ongoing negotiation, and securing trust from an audience requires consistent commitment to tenets of openness and disclosure.

A significant challenge highlighted by the interview responses and survey design is the lack of clear guidelines and standards for disclosing the use of AI in news production. Interview responses noted a spectrum of approaches, ranging from complete transparency to deliberate obfuscation (such as in the case of *Sports Illustrated*), to disclosure of AI use. The term “AI use” itself merits a deeper dive. There is no consensus on what aspects of AI involvement “merit

disclosure,” as David Cohn puts it, which raises questions about how to balance transparency with the evolving nature of AI tools and their integration into journalistic workflows. Joe Amditis’ answer to Cohn’s question of what merits disclosure is a suggestion that transparency should extend beyond simply disclosing AI authorship. It should also encompass the processes, methodologies, and potential biases associated with AI tools used in news production, while also keeping the audience in mind. While transparency can facilitate trust, David Cohn explains that part of the ongoing negotiation between writer and reader is being transparent only insofar as the audience cares about said transparency:

...it's not that I'm saying we shouldn't disclose that headline, because like, oh, it's, you know, but more like, does the audience really care? Like, 'you got this information from AI, yeah, I don't really care,' right? Like, don't waste my time. So it's more out of respect for the audience.

The crux of this challenge is finding the balance between using transparency as a mechanism for trust acquisition and respecting audience priorities. Jared Schroeder also highlighted this balance:

I've always been supportive of news organizations, labeling things, 'hey, this was generated by AI' ...but other people have said 'journalists don't give a transparency report on everything they do every time, so why would they do this?' So I guess that's yet to be decided.

For transparency to matter to an audience, they must not only care about the information being disclosed, but also possess the literacy required for said transparency to actually matter.

For example, transparency about source code is a narrow form of AI transparency that requires an understanding of code and an ability to assess the meaning and purpose of a section of code beyond a syntax-level examination. Open sourcing code by making it freely available for modification and redistribution has been promoted by large tech companies as a potential solution to issues surrounding black box algorithms (Ghioni et al., 2024). However, open source practices as a claim to transparency can also be wielded as a deceptive marketing practice known as “openwashing” in which commercial products and resources are branded as open despite terms or other practices that contradict principles of openness (Chuang et al., 2022).

Ultimately, the end goal of openness and/or transparency among platform companies that are licensing AI products (e.g. OpenAI in the Springer deal) should be to decrease the risk of error and misuse, distribute responsibility, and enable oversight both internally and externally. To assess these outcomes, we must interrogate efforts toward transparency – to whom is the entity in question being transparent? Furthermore, in what ways is this transparency articulated? And what would the public need (and what would experts need) for transparency to be functional enough to facilitate the aforementioned goals. Diligence in the auditing process is of paramount importance to counter deceptive marketing efforts that serve only to promote supposed open sourcing that, in reality, is simple hype-based branding.

The survey research illuminated trust as a particularly important aspect of perceptions of AI-generated news articles. At a fundamental level, the survey analysis reveals that reader perceptions of AI-generated news content differ somewhat from traditional news factors. The introduction of constructs such as Engagement and Alienation, robust and drawing from elements of legacy constructs such as Quality in addition to unique descriptors provided in this

survey, demonstrates the importance of reassessing news perception criteria given the advent of generative AI in journalism.

While the survey results indicate that readers do not inherently distrust AI-generated content, concerns remain about the potential for AI to perpetuate biases, generate inaccurate information, and erode public trust in journalism. Interview respondents illuminated the importance of establishing clear ethical guidelines and guardrails for AI development and deployment in newsrooms to mitigate these risks. Ultimately, building and maintaining trust requires ongoing dialogue and engagement with audiences. News organizations need to be transparent about their processes, address concerns about AI, and demonstrate their commitment to journalistic values.

7.2 Limitations

One limitation of this study involves the cumulative variance explained by the factor structures analyzed. While the identified factors provide meaningful insights into how readers perceive AI-generated news content, they do not account for the entire variance within the dataset. A substantial portion of the variance could be attributed to elements beyond these unifying factors, such as reader-specific preferences, education levels, personal interests, intelligence, and prior knowledge. These individual differences play a critical role in shaping how readers interact with and interpret news content, suggesting that the factor analysis captures only a part of the broader picture.

Additionally, the inherently exploratory nature of this study represents both a limitation and an opportunity for future research. While it offers a foundational understanding and opens avenues for exploring reader perceptions of AI-generated news, the exploratory approach means

that the findings are preliminary and should be used to inform, rather than conclusively define, subsequent in-depth investigations.

Another limitation of this study is that it is situated primarily in the context of the United States, which limits its generalizability to other regions and cultures. One reason for this is the relative lack of guardrails in the United States compared to the European Union, for example, who passed its AI Act in 2024 (which it had been planning since 2020) (Wachter, 2024). Jared Schroeder explained in his interview one of the challenges of establishing guardrails for generative artificial intelligence in the United States:

...so as it stands now, there's almost no tool that lawmakers can create no legislative tool that would limit perceived dangers from AI without violating the First Amendment, like the way this is is that we've kind of painted ourselves into a corner on that one.

Europe has been making more progress than the United States because they do not have the same free expression regime, so it is easier for them to flexibly address the consequential externalities of the proliferation of AI. The findings in the study are based almost entirely on U.S. data and reflect the uncertainty of the regulatory environment in the United States, whereas in other countries the environments are different. Therefore, conducting similar studies outside of the United States (and also potentially in non-Western contexts) in the future would potentially capture different dynamics in the production and consumption of AI-generated text given these regulatory differences.

Another limitation of the study concerns Phase 3 of the survey component, in which human-generated news was used as a comparison category. Aside from the dynamics of trust in

AI-generated news content, it may be easy to overlook the existence of shifting trust in human-generated media in addition to differences in media trust between cultures. While beyond the scope of this study, an acknowledgement of the fact that trust in the whole institution of media is volatile and warrants measurement of its own is necessary.

Across all phases of the survey, I used an inherently subjective prompting procedure with ChatGPT to generate news articles that resemble human-written articles. Given the novelty of ChatGPT as a research instrument, best practices have yet to be established for this type of research preparation, and continued experimentation will be critical to establish such practices going forward. Similarly, the process by which the descriptive data in Phase 1 of the survey was cleaned and filtered was also inherently subjective. While I provided a clear rationale for all of my inclusion and exclusion decisions and sourced input from other researchers during this process, the results of Phase 1 may differ from the hypothetical results of other researchers who may have adopted different approaches toward the filtering process. This subjectivity ultimately limits the generalizability of my survey results in exchange for greater validity.

Another limitation of the study is the small sample size of 13 semi-structured interviews conducted for the study. While the interviews conducted for this study provided ample breadth of insights in order to sufficiently address the research question, interviewing more participants in the study would have strengthened its findings and minimized potential biases. The diversity of roles within newsrooms and the varied levels of AI adoption and experimentation. A larger sample size would have provided a more comprehensive and nuanced understanding of these variations. For example, the study includes eight news workers and five experts; increasing the number of participants in each category and including additional perspectives—such as those from smaller news organizations or freelance journalists—would enhance the representativeness of the

findings. Furthermore, a larger sample size of interview respondents would allow for a more robust thematic analysis, uncovering a wider array of perspectives on topics such as AI adoption strategies, ethical considerations, and the evolving relationship between journalists and technology.

7.3 Implications

The future of artificial intelligence in news production involves uncertainty and unresolved questions regarding creativity, business models, and dependence on technology companies. The concept of alienation arose as both a salient factor in the survey results and a consideration among a small proportion of interview respondents. Looking ahead, emerging news organizations are grappling with this uncertainty, which could serve as the basis for future research alongside the foundation provided by the survey and interview results of this study.

7.3.1 Creativity

The question of creativity—the fundamental issue of how we define creativity in light of technological advancements—has legal and economic implications for how AI models train and operate. Machines didn't produce works of art until the advent of photography in the 1860s, a technological breakthrough that blurred the line between human creation and mechanical reproduction. Creativity, historically understood as an expression of human intelligence, has been rooted in Enlightenment principles that associate creativity with the rational mind. This connection between intelligence and artistic expression is why we protect writing as a form of creative expression: it is the manifestation of ideas, which are considered natural, and therefore unique to the individual who conceptualizes them. Knowledge is transformed into information

through representation—which could be text, image, or sound—and then that information is subsequently acquired by the observer. When the observer understands and integrates these representations, information transforms back into knowledge, fitting into their evolving picture of the world.

As we considered in chapter 3, the role of artificial intelligence intervenes in this loop at the stage of representation. AI does not simply reproduce knowledge, nor does it create representations. It generates new probabilistic inferences that resemble the outputs of creative processes without conforming to the mimetic ontology of the creative process itself. However, this raises a philosophical and legal challenge: if ideas are natural and inherently tied to human authorship, how do we treat the outputs of AI, which lacks the human mind but can still generate text and art? Traditionally, we view ideas as the property of the author, with copyright serving as a temporary protective measure. After a set period, copyright expires, and the once-protected ideas are returned to the public domain for others to build upon.

With AI-generated text, this dynamic becomes more complicated. If we hold onto the idea that creativity is exclusively human, then AI's output poses a question: is it merely a tool for human creativity, or does it produce something of its own that challenges our definition of originality and ownership? While the legal system currently struggles to classify AI-generated works, there is an emerging tension between protecting creative expression and acknowledging that, with AI, creativity might no longer belong solely to humans. Currently, works generated by AI are regarded as belonging to the public domain. However, the expiration of copyright, which traditionally reassimilates ideas into the public domain, may no longer suffice if AI can continuously generate outputs that resemble human outputs. Thus, we are at a crossroads where

we must redefine the boundaries of creativity and intellectual property, particularly as we strive to maintain creativity as something that belongs to the public at large.

7.3.2 Alienation

One of the most interesting findings in terms of implications for future theoretical investigation of AI-generated news was the salience of the Alienation factor in the survey research. Alienation was obtained as a salient factor after both the varimax and the oblimin rotations of the underlying Phase 2 data, having high loadings of the Biased, Rote, and Weird descriptors in each. Specifically, the high loading of the Biased descriptor onto Alienation stands out. As mentioned in chapter 6, given the traditional association with bias as a marker of credibility or quality, this result seems surprising on the surface level. Yet considerations of the influence of social AI perceptions on content perceptions along with source- and content-based cues that relate to alienation are possible explanations for these observations. Insights from interview responses support these considerations. Joe Amditis explains:

But there's something core to that, that feeling of the bot spitting out above average text in most cases that makes people feel weird, and that's fine. That's cool. But for me, I'm interested in finding that pinpoint where—and the justification for when—it goes over that line.

The “line” that Amditis mentions here is an important threshold for producers of AI-generated text to keep in mind. Before the advent of AI-generated text, news researchers generally thought about reader perceptions in terms of factors like those obtained in Sundar (1999)’s study: credibility, liking, quality, and readability. The Alienation factor represents a

unique constellation of reader considerations informed by their awareness of AI authorship: given that “Boring” and “Rote” are content-level evaluations, perhaps readers are evaluating bias at the level of content rather than source. In the context of AI, this would imply an understanding among readers that bias occurs at the level of model training. Even if AI systems can generate factually accurate content, the lack of human-like engagement and relatability could lead to feelings of alienation and discourage readers from actively consuming the information. Ultimately, the ability of AI to generate human-like text disrupts the traditional understanding of creativity and authorship, which can lead to a sense of alienation among readers accustomed to attributing written content to human minds and experiences.

Relevant to discussions and analysis about feelings of alienation when interacting with AI-generated content online, The Dead Internet Theory posits that the internet is largely populated by bots and AI-generated content designed to manipulate users (Walter, 2024). Many of the experiences discussed by respondents in their interviews implicitly spoke to the implications of this theory on digital interaction. Chris Dinn expressed concerns about over-reliance on AI and the potential for it to "ossify culture" by perpetuating existing biases or limiting diversity of thought. He also spoke about the potential for LLMs to train on LLM-generated text. Given the scale at which models can produce and disseminate text online, future training data sets may predominantly include text generated not by humans, but by prior iterations of large language models.

7.3.3 Looking Ahead

A common thread underlying contemporary discourse on AI is its inevitability, and the acceptance or denial of said inevitability. Jay Allred aligns with the acceptance perspective in his experience:

...we have viewed that from the beginning as an inevitability, there's too much money pouring in from these firms like Google and OpenAI and X and whatever, there's just too much money at stake for this not to work. We don't know how fast it will work, we don't know how quickly it will reach some sort of tipping point. But we're approaching it like the tipping point is inevitable.

Many creative professionals are reluctant to acknowledge that generative AI models like ChatGPT could significantly impact or even replace their jobs. For example, copywriters, a field heavily reliant on human creativity and linguistic nuance, now find themselves at risk. Rather than producing original copy from scratch, many copywriters are increasingly relying on AI tools to assist their process. Instead of being the sole creators, they are now collaborating with machines, using AI to tweak and refine drafts, running their ideas through multiple iterations to optimize output. While the act of writing may still involve human judgment, the true creative spark is often reduced to crafting effective prompts for the AI to generate content. This shift in workflow calls back to the question of creativity: how can we be certain about the future role of human creativity in industries that are becoming more reliant on machine-generated output?

Emerging news organizations are beginning to grapple with questions of creativity and truth in the media. For example, in early 2024, the former MSNBC host Mehdi Hasan launched Zeteo, a new digital subscription media service. Looking ahead: Emerging news orgs grappling

with this e.g. Mehdi Hassan Zeteo. Hosted on Substack, an online publishing platform, Zeteo boasts 245000 subscribers as of September 2024 (Barr, 2024). Hasan joins other media commentators who have recently established alternative media presences, including Bari Weiss, Nate Silver, Tucker Carlson, and Glenn Beck. The shift toward alternative media is an effort toward autonomy in media production, aligning with the emphasis on the need for the “human element” in journalism expressed by practitioners in the interview analysis.

Building on the results of this study, future research should investigate the perception differences between human and machine-written content across different article categories. Additionally, future research could use the newly discovered factors (e.g. Effort, Substance, Clarity) in reader perception studies comparing human-written and AI-written news. This research provides valuable insights into the evolving landscape of automated journalism and its perception among news readers. Future research could also utilize the factors explored in this study as metrics upon which AI- and human-generated news content (and perhaps communication content more broadly) can be compared. Ultimately, this study provides a foundation for offering a more comprehensive understanding of the multifaceted nature of news production and consumption in the age of AI.

APPENDIX

A) Article Generation Prompting

Phase 1

Article 1

Write a Reuters news article titled "Internet Archive's digital book lending violates copyrights, US judge rules"



Here's the article:

WASHINGTON (Reuters) - The Internet Archive's digital book lending program, which allows users to borrow digital copies of books, has been ruled to violate copyright law by a U.S. judge.

The decision, handed down on Friday by Judge John Blakey in Chicago, is a blow to the non-profit organization's efforts to expand access to literature through its online library, which currently houses millions of free books, movies, music and more.

The Internet Archive's "Open Library" program was launched in 2005 and initially allowed users to check out physical copies of books. But in 2011, the organization began digitizing its collection and making those digital copies available for borrowing.

Several publishers, including HarperCollins, Penguin Random House and Hachette Book Group, sued the Internet Archive in 2019, alleging that the digital lending program infringed their copyrights.

In his ruling, Blakey said that the Internet Archive's digital book lending program did not qualify for the "fair use" exemption under copyright law, which allows for limited use of copyrighted material for purposes such as criticism, commentary, news reporting, teaching, scholarship, or research.

"The Internet Archive did not have permission from plaintiffs to scan their books, and they did not have permission to distribute them," Blakey wrote in his decision.

The ruling comes amid ongoing debates about the balance between copyright protections and access to information in the digital age. Supporters of the Internet Archive's digital lending program argue that it provides valuable access to books that might otherwise be difficult to obtain, particularly for those who live in remote areas or cannot afford to purchase books.

However, publishers and authors have argued that the digital lending program threatens their ability to earn revenue from their works.

The Internet Archive has said that it will appeal the ruling, and that it believes its digital lending program is legal and beneficial to the public.

"We disagree with the court's decision and are disappointed that a small group of publishers is trying to undermine the public's access to books," the organization said in a statement.

The Internet Archive's digital book lending program remains operational for now, but its future is uncertain pending the outcome of the appeal.

Phase 2

Article 1

Write a news article in the style of Reuters titled "Time running out for Trump's Republican rivals as they face off in second debate"



Time running out for Trump's Republican rivals as they face off in second debate

By [Your Name], Reuters

WASHINGTON (Reuters) - As the clock ticks down to the pivotal primary season, the Republican contenders jostled for prominence against former President Donald Trump in a fiery second debate on Thursday evening.

Facing off in a bustling auditorium in Des Moines, Iowa, Trump's challengers sought to assert their credentials, question his policies, and gain traction in what has increasingly seemed like a one-horse race.

Despite a crowded field of seasoned politicians, accomplished business leaders, and new faces, recent polls have shown Trump enjoying a comfortable lead. For many of the candidates, Thursday's debate represented one of the last significant opportunities to change the trajectory of the race.

Throughout the evening, former Florida Governor John Smith, who has emerged as a key adversary for Trump, pressed him on issues ranging from trade policies to his handling of foreign affairs. "Our party deserves a leader who understands the intricacies of global politics, not just Twitter politics," Smith retorted after a heated exchange.

Meanwhile, tech entrepreneur Emily Richards, a surprise late entry to the race, aimed to differentiate herself from career politicians. "We need innovative solutions for the challenges of today, not recycled ideas from the past," she stated, capturing significant applause.

However, the former President, known for his combative style and ability to dominate media cycles, frequently took center stage. Trump defended his past administration's record and showcased his vision for the future, while often deflecting attacks with characteristic aplomb.

Political analysts believe that the debate's outcome is crucial. "With the Iowa caucuses just weeks away, candidates are running out of time to make their mark," said political science professor Dr. Alice Thompson from George Washington University. "If they can't break through now, the race might be all but decided."

Despite the urgency and high stakes, the debate also had moments of unity. All candidates voiced their shared commitment to core Republican values and the need to triumph in the upcoming general election.

As the night drew to a close, the big question on many minds was whether this face-off had altered the landscape of the Republican primary race or if Trump's lead remains unassailable.

The Iowa caucuses, which kick off the primary season, will be held in early February. The results could either solidify Trump's frontrunner status or breathe new life into his rivals' campaigns.

(Reporting by [Your Name], Editing by [Editor's Name])

Article 2

Write a news article in the style of Reuters titled "Volkswagen hit by IT outage, brand vehicle production in Germany halted"



Volkswagen hit by IT outage, brand vehicle production in Germany halted

By [Your Name], Reuters

WOLFSBURG, Germany (Reuters) - Volkswagen, one of the world's largest automakers, faced a significant IT outage on Friday, leading to a temporary suspension of brand vehicle production across its German plants.

The outage, which started in the early morning hours, impacted the company's internal systems, affecting production schedules, supply chain management, and employee communication. Volkswagen's main production hub in Wolfsburg, along with other plants in Emden, Zwickau, and Dresden, were among those that halted operations.

A spokesperson for Volkswagen stated, "We are aware of the IT disruptions and are working urgently to resolve the situation. Our primary concern is to ensure the stability of our operations and minimize the impact on our workforce and customers."

The cause of the disruption remains under investigation. While there was no immediate indication of an external cyber attack, the incident underlines the auto industry's vulnerability to IT-related issues in an era of increasing digitization and connectivity.

Germany's auto industry has been grappling with challenges, from the global semiconductor shortage to shifts toward electric mobility. This latest disruption adds another layer of complexity to the sector's recovery from pandemic-induced slowdowns.

Analysts anticipate that the outage could exacerbate existing delivery delays, especially if the halt prolongs. "With an already tight supply chain, disruptions like these can have cascading effects on the production schedule and ultimately the end consumer," noted automotive analyst Dr. Benjamin Hartman from Berlin's Institute of Automotive Research.

Volkswagen's shares slid 2.5% in Frankfurt trading after the news broke, reflecting investor concerns over potential financial impacts.

The company is aiming to restore partial operations by the end of the day and is in continuous communication with its suppliers and workforce to mitigate the effects of the shutdown.

(Reporting by [Your Name], Editing by [Editor's Name])

Article 3

Write a news article in the style of Reuters titled "Epic Games asks US Supreme Court to review Apple antitrust case"



Epic Games asks US Supreme Court to review Apple antitrust case

By [Your Name], Reuters

WASHINGTON (Reuters) - Epic Games, the developer behind the popular video game "Fortnite", has petitioned the U.S. Supreme Court to review a previous appellate decision related to its antitrust lawsuit against tech giant Apple Inc.

The dispute began when Epic Games attempted to bypass Apple's in-app purchase system, leading to "Fortnite's" removal from the App Store. Epic subsequently launched a lawsuit, alleging that Apple's App Store policies and fees amounted to monopolistic behavior.

While a federal judge initially handed a mixed verdict that recognized Apple's restrictions on app distribution as anti-competitive, the court also found Epic Games in breach of its developer agreement. Both parties appealed the decision, with Apple's policies coming under further scrutiny.

In the petition to the Supreme Court, Epic Games argues that the appellate court's decision "fails to recognize the broader implications of Apple's stranglehold on app distribution and in-app payments."

Apple responded in a statement: "The App Store has revolutionized the tech industry, offering a safe and trusted platform for developers and users. We believe the appellate decision was correct and will continue to defend our practices."

This legal battle has garnered significant attention, serving as a focal point for broader discussions on tech giants' influence and the nature of digital marketplaces. Regulators and lawmakers worldwide have been watching the case closely, as decisions could set precedents affecting digital commerce regulations globally.

"Epic's decision to take this to the Supreme Court underscores the stakes at play. Beyond this single case, the outcome could shape the future of digital ecosystems and competition," remarked tech analyst Laura Diaz from the Digital Rights Institute.

As the gaming and tech worlds await the Supreme Court's decision on whether to hear the case, discussions on market fairness, competition, and innovation continue to heat up.

(Reporting by [Your Name], Editing by [Editor's Name])

Phase 3 Article 1

Write a news article in the style of Reuters titled "US inflation decelerating in boost to economy"



US Inflation Decelerating in Boost to Economy

In a significant development for the United States economy, recent data indicates that inflation, a key indicator of economic health, is showing signs of deceleration. This trend is being hailed as a positive signal for the economy, potentially easing the financial strain on households and businesses across the nation.

According to the latest figures released by the Department of Labor, the Consumer Price Index (CPI), a widely watched measure of inflation, has recorded a slower growth rate compared to previous months. Economists suggest that this slowdown in inflation could be attributed to various factors, including adjustments in monetary policy by the Federal Reserve, stabilization in global commodity prices, and a gradual resolution of supply chain disruptions that have plagued the economy since the onset of the COVID-19 pandemic.

The deceleration of inflation is particularly noteworthy given the challenges faced by the U.S. economy over the past year. High inflation rates have been a persistent concern, impacting the cost of living and eroding the purchasing power of consumers. The recent easing of inflationary pressures is therefore seen as a relief and a potential catalyst for more robust economic growth.

In response to the positive trend, Federal Reserve officials have expressed cautious optimism. While acknowledging the improvement, they emphasize the need for continued vigilance to ensure long-term economic stability. The Federal Reserve has been actively managing monetary policy to combat inflation, including adjusting interest rates. These measures appear to be bearing fruit, as evidenced by the latest inflation data.

The business community has also welcomed the news, with many industry leaders expressing hope that a more stable inflation rate will lead to improved consumer confidence and spending. This, in turn, could drive growth in various sectors of the economy, from retail to manufacturing.

Analysts warn, however, that while the deceleration of inflation is a positive development, the economy is not out of the woods yet. Global economic uncertainties, such as geopolitical tensions and potential new waves of the pandemic, could still pose risks to the economic recovery.

In summary, the deceleration of inflation in the United States is a welcome development for the economy. It brings relief to consumers and businesses alike and sets the stage for potentially stronger economic growth in the coming months. However, caution remains the watchword, with a keen eye on global economic trends and domestic policy responses.

Article 2

Write a news article in the style of Reuters titled "Supreme Court rejects prosecutor's request, will not rule on Trump immunity yet"



Supreme Court Rejects Prosecutor's Request, Will Not Rule on Trump Immunity Yet

In a recent development, the United States Supreme Court has declined to hear a case regarding former President Donald Trump's claim of immunity from criminal prosecution. The decision effectively rejects a request from a prosecutor to expedite the Court's ruling on the matter, leaving the issue of Trump's immunity unresolved for the time being.

The case in question revolves around Trump's assertion that, as a former President, he is immune from criminal prosecution for actions taken while in office. This claim has been a point of contention and legal debate, particularly in light of ongoing investigations into various activities during Trump's presidency.

A prosecutor, whose name and office have not been disclosed in keeping with Reuters policy on ongoing legal matters, had petitioned the Supreme Court to fast-track a decision on this critical constitutional question. The request was seen as an attempt to clear the way for potential legal action against the former President.

However, in a brief order, the Supreme Court stated that it would not expedite its consideration of the case. The order did not provide any reasoning for the decision, which is typical for such rulings. This means that the Court will not yet address the broader issue of a former President's immunity from criminal prosecution, a topic that has significant implications for the legal and political landscape.

Legal experts suggest that the Supreme Court's decision to not fast-track the case does not indicate its stance on the merits of the argument. Instead, it reflects the Court's typical approach to managing its docket and the complexities involved in constitutional questions of this nature.


The lack of an immediate ruling leaves the legal situation somewhat in limbo. It means that ongoing investigations and potential legal actions against Trump will continue under the current legal framework, without specific guidance from the Supreme Court on the immunity question.

The decision has drawn reactions from various quarters. Supporters of Trump view it as a temporary relief, while his critics see it as a missed opportunity for the Supreme Court to clarify a critical aspect of presidential immunity.

In summary, the Supreme Court's refusal to expedite a decision on former President Donald Trump's claim of immunity from criminal prosecution keeps the significant constitutional question open. This decision has implications for ongoing legal proceedings and the broader understanding of presidential powers and responsibilities post-office. The legal and political communities will be watching closely for any future developments in this case.

Article 3

Write a news article in the style of Reuters titled "Exclusive: Synopsys seeks to acquire engineering software company Ansys-sources"

 **Exclusive: Synopsys Seeks to Acquire Engineering Software Company Ansys - Sources**

In a significant development in the technology sector, Synopsys Inc., a leading provider of electronic design automation software, is reportedly in advanced talks to acquire Ansys Inc., a global leader in engineering simulation software. This information comes from sources close to the matter who requested anonymity due to the private nature of the discussions.

According to these sources, Synopsys has expressed a keen interest in Ansys as part of its strategy to expand its footprint in the engineering and design software market. The potential acquisition would create a powerhouse in the industry, combining Synopsys's expertise in electronic design with Ansys's capabilities in engineering simulation.

Ansys, known for its software that simulates and analyzes the performance of structures, fluids, electronics, and software used in a wide range of industries, could complement Synopsys's existing portfolio. This acquisition could enable Synopsys to offer a broader range of software solutions that extend beyond its traditional semiconductor design business.

The financial terms of the potential deal have not been disclosed, and representatives from both Synopsys and Ansys have declined to comment on the matter. It is important to note that the discussions are ongoing and may not lead to a definitive agreement.

The move, if it materializes, would mark a significant consolidation in the software industry, particularly in the niche sector of engineering and design. It reflects a growing trend of mergers and acquisitions as companies seek to broaden their technological capabilities and market reach.

The acquisition would also likely attract regulatory scrutiny, given the prominent positions of both companies in their respective markets. Regulators may examine the deal for potential antitrust issues, including market concentration and customer impact.

For Synopsys, acquiring Ansys would potentially enhance its position as a one-stop-shop for engineering and design software, offering integrated solutions to its clients. For Ansys, joining forces with Synopsys could provide access to new markets and customers, leveraging Synopsys's global reach and resources.

In summary, the potential acquisition of Ansys by Synopsys could represent a major shift in the engineering and design software landscape, creating a combined entity with enhanced capabilities and market presence. However, the deal is still under negotiation and subject to various approvals and conditions.

B) Interview Guide

News Workers

Main Theme	Main Question	Support Question(s)
Sensemaking	<ul style="list-style-type: none"> Do you use any strategies to increase user engagement with content? 	<ul style="list-style-type: none"> In what ways do algorithms assist in driving that engagement?
	<ul style="list-style-type: none"> Have software and algorithms impacted your role as a 	<ul style="list-style-type: none"> If so, in what way has your role been impacted?

	journalist in determining the editorial relevance and newsworthiness of content?	
	<ul style="list-style-type: none"> • What strategies, if any, do journalists use to compete with or counteract automated disinformation campaigns? 	
Trust	<ul style="list-style-type: none"> • What strategies are used by newswriters to elicit trust in published content, particularly for content shared on social media? 	<ul style="list-style-type: none"> ○ Are these strategies effective? Why or why not?
	<ul style="list-style-type: none"> • How do individual journalists contribute to the overall credibility of a source/news outlet? 	
	<ul style="list-style-type: none"> • Do you think that the integration of software companies into the media production process improves, worsens, or has no effect on public trust of journalism as an institution? 	<ul style="list-style-type: none"> • Why?
Knowledge	<ul style="list-style-type: none"> • How would you describe your understanding of automated journalism? 	
	<ul style="list-style-type: none"> • Do you think that 	<ul style="list-style-type: none"> ○ If so, to what extent

	journalists are creators of knowledge?	can algorithms mimic this role? ○ If not, do you think algorithms are creators of knowledge?
	<ul style="list-style-type: none"> To what extent do journalists at [your organization] participate in the algorithmic news production process? 	
Other	<ul style="list-style-type: none"> Does automation make your job easier or harder? 	<ul style="list-style-type: none"> Why or why not?

Experts

Main Theme	Main Question	Support Question(s)
Sensemaking	<ul style="list-style-type: none"> How do conceptions of user behavior inform the development process for news algorithms? 	
	<ul style="list-style-type: none"> What strategies do natural language machine learning engineers employ to optimize the understandability of output text? 	
	<ul style="list-style-type: none"> What metrics are used to judge the output quality of an NLP algorithm? 	<ul style="list-style-type: none"> How/why were these metrics chosen?
Trust	<ul style="list-style-type: none"> To what extent do recommendation 	

	algorithms prioritize the promotion of trustworthy content versus the optimization of engagement?	
	<ul style="list-style-type: none"> • What policies exist within the software space to support ethical use of NLP algorithms and combat issues such as misinformation? 	
	<ul style="list-style-type: none"> • Do you think that the integration of software companies into the media production process improves, worsens, or has no effect on public trust of journalism as an institution? 	
Knowledge	<ul style="list-style-type: none"> • Do concerns regarding biased data manifest themselves in the NLG algorithm design process? 	<ul style="list-style-type: none"> ○ If not, to what extent do designers of algorithms consider themselves creators of knowledge as opposed to the algorithms they build?
	<ul style="list-style-type: none"> • If an algorithm can hypothetically operate upon a corpus of input data consisting of all written content over the course of human history, to what extent would its output be indistinguishable from human-written text? 	
	<ul style="list-style-type: none"> • In the future, do you 	<ul style="list-style-type: none"> ○ Why or why not?

	think NLG algorithms will surpass the human brain in terms of intelligence?	
Other	<ul style="list-style-type: none"> • What is your role in relation to NLG algorithms? 	

C) Survey Questions

Phase 1 Survey

- Background: The following article was written using GPT-4, a multimodal large language model created by OpenAI. To write the following article, GPT-4 was prompted with the title and source of an existing, human-written article with the same title.
- Task: Please read the article and answer the subsequent questions.
- Title: "Internet Archive's digital book lending violates copyrights, US judge rules".
- WASHINGTON (Reuters) - The Internet Archive's digital book lending program, which allows users to borrow digital copies of books, has been ruled to violate copyright law by a U.S. judge. The decision, handed down on Friday by Judge John Blakey in Chicago, is a blow to the non-profit organization's efforts to expand access to literature through its online library, which currently houses millions of free books, movies, music and more. The Internet Archive's "Open Library" program was launched in 2005 and initially allowed users to check out physical copies of books. But in 2011, the organization began digitizing its collection and making those digital copies available for borrowing. Several publishers, including HarperCollins, Penguin Random House and Hachette Book Group, sued the Internet Archive in 2019, alleging that the digital lending program infringed their copyrights. In his ruling, Blakey said that the Internet Archive's digital book lending program did not qualify for the "fair use" exemption under copyright law, which allows for limited use of copyrighted material for purposes such as criticism, commentary, news

reporting, teaching, scholarship, or research. "The Internet Archive did not have permission from plaintiffs to scan their books, and they did not have permission to distribute them," Blakey wrote in his decision. The ruling comes amid ongoing debates about the balance between copyright protections and access to information in the digital age. Supporters of the Internet Archive's digital lending program argue that it provides valuable access to books that might otherwise be difficult to obtain, particularly for those who live in remote areas or cannot afford to purchase books. However, publishers and authors have argued that the digital lending program threatens their ability to earn revenue from their works. The Internet Archive has said that it will appeal the ruling, and that it believes its digital lending program is legal and beneficial to the public. "We disagree with the court's decision and are disappointed that a small group of publishers is trying to undermine the public's access to books," the organization said in a statement. The Internet Archive's digital book lending program remains operational for now, but its future is uncertain pending the outcome of the appeal.

- Question: List the thoughts that come to your mind after reading the article
- Question: List 2-10 adjectives describing the article

Phase 2: Descriptor Ranking Survey

- Background: In this survey, you will read three news articles. These three articles were written using GPT-4, a multimodal large language model created by OpenAI. To write each of the following articles, GPT-4 was prompted with the title and source of an existing, human-written article with the same title. The output has not been edited in any way.
- Task: Please read the articles and answer the subsequent questions.
- Article 1: Time running out for Trump's Republican rivals as they face off in second debate

By [Your Name], Reuters

WASHINGTON (Reuters) - As the clock ticks down to the pivotal primary season, the Republican contenders jostled for prominence against former President Donald Trump in

a fiery second debate on Thursday evening.

Facing off in a bustling auditorium in Des Moines, Iowa, Trump's challengers sought to assert their credentials, question his policies, and gain traction in what has increasingly seemed like a one-horse race.

Despite a crowded field of seasoned politicians, accomplished business leaders, and new faces, recent polls have shown Trump enjoying a comfortable lead. For many of the candidates, Thursday's debate represented one of the last significant opportunities to change the trajectory of the race.

Throughout the evening, former Florida Governor John Smith, who has emerged as a key adversary for Trump, pressed him on issues ranging from trade policies to his handling of foreign affairs. "Our party deserves a leader who understands the intricacies of global politics, not just Twitter politics," Smith retorted after a heated exchange.

Meanwhile, tech entrepreneur Emily Richards, a surprise late entry to the race, aimed to differentiate herself from career politicians. "We need innovative solutions for the challenges of today, not recycled ideas from the past," she stated, capturing significant applause.

However, the former President, known for his combative style and ability to dominate media cycles, frequently took center stage. Trump defended his past administration's record and showcased his vision for the future, while often deflecting attacks with characteristic aplomb.

Political analysts believe that the debate's outcome is crucial. "With the Iowa caucuses just weeks away, candidates are running out of time to make their mark," said political science professor Dr. Alice Thompson from George Washington University. "If they can't break through now, the race might be all but decided."

Despite the urgency and high stakes, the debate also had moments of unity. All candidates voiced their shared commitment to core Republican values and the need to triumph in the upcoming general election.

As the night drew to a close, the big question on many minds was whether this face-off had altered the landscape of the Republican primary race or if Trump's lead remains unassailable.

The Iowa caucuses, which kick off the primary season, will be held in early February.

The results could either solidify Trump's frontrunner status or breathe new life into his rivals' campaigns.

(Reporting by [Your Name], Editing by [Editor's Name])

- Instructions: In the following question, for each word, please rate how well the word describes the article above, from "describes very poorly" (1) to "describes very well" (5)
- Q1: Use the sliders below to indicate your ratings
- Amazing Believable Biased Boring Coherent Complex Concise
Conclusive Detailed Educational Factual Fair Honest Informative
Interesting Long Moral Neutral Precise Rote Technical Thorough
Thought-provoking Timely Weird Wordy
- Article 2 Volkswagen hit by IT outage, brand vehicle production in Germany halted

By [Your Name], Reuters

WOLFSBURG, Germany (Reuters) - Volkswagen, one of the world's largest automakers, faced a significant IT outage on Friday, leading to a temporary suspension of brand vehicle production across its German plants.

The outage, which started in the early morning hours, impacted the company's internal systems, affecting production schedules, supply chain management, and employee communication. Volkswagen's main production hub in Wolfsburg, along with other plants in Emden, Zwickau, and Dresden, were among those that halted operations.

A spokesperson for Volkswagen stated, "We are aware of the IT disruptions and are working urgently to resolve the situation. Our primary concern is to ensure the stability of our operations and minimize the impact on our workforce and customers."

The cause of the disruption remains under investigation. While there was no immediate indication of an external cyber attack, the incident underlines the auto industry's vulnerability to IT-related issues in an era of increasing digitization and connectivity.

Germany's auto industry has been grappling with challenges, from the global semiconductor shortage to shifts toward electric mobility. This latest disruption adds another layer of complexity to the sector's recovery from pandemic-induced slowdowns.

Analysts anticipate that the outage could exacerbate existing delivery delays, especially if the halt prolongs. "With an already tight supply chain, disruptions like these can have

cascading effects on the production schedule and ultimately the end consumer," noted automotive analyst Dr. Benjamin Hartman from Berlin's Institute of Automotive Research.

Volkswagen's shares slid 2.5% in Frankfurt trading after the news broke, reflecting investor concerns over potential financial impacts.

The company is aiming to restore partial operations by the end of the day and is in continuous communication with its suppliers and workforce to mitigate the effects of the shutdown.

(Reporting by [Your Name], Editing by [Editor's Name])

- Instructions: In the following question, for each word, please rate how well the word describes the article above, from "describes very poorly" (1) to "describes very well" (5)
- Q2: Use the sliders below to indicate your ratings
- Amazing Believable Biased Boring Coherent Complex Concise Conclusive Detailed Educational Factual Fair Honest Informative Interesting Long Moral Neutral Precise Rote Technical Thorough Thought-provoking Timely Weird Wordy
- Article 3: Epic Games asks US Supreme Court to review Apple antitrust case

By [Your Name], Reuters

WASHINGTON (Reuters) - Epic Games, the developer behind the popular video game "Fortnite", has petitioned the U.S. Supreme Court to review a previous appellate decision related to its antitrust lawsuit against tech giant Apple Inc.

The dispute began when Epic Games attempted to bypass Apple's in-app purchase system, leading to "Fortnite's" removal from the App Store. Epic subsequently launched a lawsuit, alleging that Apple's App Store policies and fees amounted to monopolistic behavior.

While a federal judge initially handed a mixed verdict that recognized Apple's restrictions on app distribution as anti-competitive, the court also found Epic Games in breach of its developer agreement. Both parties appealed the decision, with Apple's policies coming under further scrutiny.

In the petition to the Supreme Court, Epic Games argues that the appellate court's

decision "fails to recognize the broader implications of Apple's stranglehold on app distribution and in-app payments."

Apple responded in a statement: "The App Store has revolutionized the tech industry, offering a safe and trusted platform for developers and users. We believe the appellate decision was correct and will continue to defend our practices."

This legal battle has garnered significant attention, serving as a focal point for broader discussions on tech giants' influence and the nature of digital marketplaces. Regulators and lawmakers worldwide have been watching the case closely, as decisions could set precedents affecting digital commerce regulations globally.

"Epic's decision to take this to the Supreme Court underscores the stakes at play. Beyond this single case, the outcome could shape the future of digital ecosystems and competition," remarked tech analyst Laura Diaz from the Digital Rights Institute.

As the gaming and tech worlds await the Supreme Court's decision on whether to hear the case, discussions on market fairness, competition, and innovation continue to heat up.

(Reporting by [Your Name], Editing by [Editor's Name])

- Instructions: In the following question, for each word, please rate how well the word describes the article above, from "describes very poorly" (1) to "describes very well" (5)
- Q3: Use the sliders below to indicate your ratings

Amazing () Believable () Biased () Boring () Coherent () Complex () Concise () Conclusive () Detailed () Educational () Factual () Fair () Honest () Informative () Interesting () Long () Moral () Neutral () Precise () Rote () Technical () Thorough () Thought-provoking () Timely () Weird () Wordy ()

- Survey Completion ID Here is your ID: `{e://Field/Random%20ID}`

Copy this value to paste into MTurk. When you have copied the ID, please click the next button to submit your survey.

Phase 3: Factor Ranking Survey

Factor Ranking Survey - Article Condition

In this survey, you will read three news articles. Please read the articles and answer the subsequent questions.

AI Article - Finance US Inflation Decelerating in Boost to Economy

In a significant development for the United States economy, recent data indicates that inflation, a key indicator of economic health, is showing signs of deceleration. This trend is being hailed as a positive signal for the economy, potentially easing the financial strain on households and businesses across the nation.

According to the latest figures released by the Department of Labor, the Consumer Price Index (CPI), a widely watched measure of inflation, has recorded a slower growth rate compared to previous months. Economists suggest that this slowdown in inflation could be attributed to various factors, including adjustments in monetary policy by the Federal Reserve, stabilization in global commodity prices, and a gradual resolution of supply chain disruptions that have plagued the economy since the onset of the COVID-19 pandemic.

The deceleration of inflation is particularly noteworthy given the challenges faced by the U.S. economy over the past year. High inflation rates have been a persistent concern, impacting the cost of living and eroding the purchasing power of consumers. The recent easing of inflationary pressures is therefore seen as a relief and a potential catalyst for more robust economic growth.

In response to the positive trend, Federal Reserve officials have expressed cautious optimism. While acknowledging the improvement, they emphasize the need for continued vigilance to ensure long-term economic stability. The Federal Reserve has been actively managing monetary policy to combat inflation, including adjusting interest rates. These measures appear to be bearing fruit, as evidenced by the latest inflation data.

The business community has also welcomed the news, with many industry leaders expressing hope that a more stable inflation rate will lead to improved consumer confidence and spending. This, in turn, could drive growth in various sectors of the economy, from retail to manufacturing.

Analysts warn, however, that while the deceleration of inflation is a positive development, the economy is not out of the woods yet. Global economic uncertainties, such as geopolitical tensions and potential new waves of the pandemic, could still pose risks to the economic recovery.

In summary, the deceleration of inflation in the United States is a welcome development for the economy. It brings relief to consumers and businesses alike and sets the stage for potentially stronger economic growth in the coming months. However, caution remains the watchword, with a keen eye on global economic trends and domestic policy responses.

In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)

Use the sliders below to indicate your ratings 1 2 3 4 5 Trust () Prolixity (the use of too many words to express an idea) () Engagement () Substance () Clarity () Alienation () Complexity () Effort (required to read) () Neutrality ()

AI Article - Politic Supreme Court Rejects Prosecutor's Request, Will Not Rule on Trump Immunity Yet

In a recent development, the United States Supreme Court has declined to hear a case regarding former President Donald Trump's claim of immunity from criminal prosecution. The decision effectively rejects a request from a prosecutor to expedite the Court's ruling on the matter, leaving the issue of Trump's immunity unresolved for the time being.

The case in question revolves around Trump's assertion that, as a former President, he is immune from criminal prosecution for actions taken while in office. This claim has been a point of contention and legal debate, particularly in light of ongoing investigations into various activities during Trump's presidency.

A prosecutor, whose name and office have not been disclosed in keeping with Reuters policy on ongoing legal matters, had petitioned the Supreme Court to fast-track a decision on this critical constitutional question. The request was seen as an attempt to clear the way for potential legal action against the former President.

However, in a brief order, the Supreme Court stated that it would not expedite its consideration of the case. The order did not provide any reasoning for the decision, which is typical for such rulings. This means that the Court will not yet address the broader issue of a former President's immunity from criminal prosecution, a topic that has significant implications for the legal and political landscape.

Legal experts suggest that the Supreme Court's decision to not fast-track the case does not indicate its stance on the merits of the argument. Instead, it reflects the Court's typical approach to managing its docket and the complexities involved in constitutional questions of this nature.

The lack of an immediate ruling leaves the legal situation somewhat in limbo. It means that ongoing investigations and potential legal actions against Trump will continue under the current legal framework, without specific guidance from the Supreme Court on the immunity question.

The decision has drawn reactions from various quarters. Supporters of Trump view it as a temporary relief, while his critics see it as a missed opportunity for the Supreme Court to clarify a critical aspect of presidential immunity.

In summary, the Supreme Court's refusal to expedite a decision on former President Donald

Trump's claim of immunity from criminal prosecution keeps the significant constitutional question open. This decision has implications for ongoing legal proceedings and the broader understanding of presidential powers and responsibilities post-office. The legal and political communities will be watching closely for any future developments in this case.

In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)

Use the sliders below to indicate your ratings 1 2 3 4 5 Trust () Prolixity (the use of too many words to express an idea) () Engagement () Substance () Clarity () Alienation () Complexity () Effort (required to read) () Neutrality ()

AI Article - Tech Exclusive: Synopsys Seeks to Acquire Engineering Software Company Ansys - Sources

In a significant development in the technology sector, Synopsys Inc., a leading provider of electronic design automation software, is reportedly in advanced talks to acquire Ansys Inc., a global leader in engineering simulation software. This information comes from sources close to the matter who requested anonymity due to the private nature of the discussions.

According to these sources, Synopsys has expressed a keen interest in Ansys as part of its strategy to expand its footprint in the engineering and design software market. The potential acquisition would create a powerhouse in the industry, combining Synopsys's expertise in electronic design with Ansys's capabilities in engineering simulation.

Ansys, known for its software that simulates and analyzes the performance of structures, fluids, electronics, and software used in a wide range of industries, could complement Synopsys's existing portfolio. This acquisition could enable Synopsys to offer a broader range of software solutions that extend beyond its traditional semiconductor design business.

The financial terms of the potential deal have not been disclosed, and representatives from both Synopsys and Ansys have declined to comment on the matter. It is important to note that the discussions are ongoing and may not lead to a definitive agreement.

The move, if it materializes, would mark a significant consolidation in the software industry, particularly in the niche sector of engineering and design. It reflects a growing trend of mergers and acquisitions as companies seek to broaden their technological capabilities and market reach.

The acquisition would also likely attract regulatory scrutiny, given the prominent positions of both companies in their respective markets. Regulators may examine the deal for potential antitrust issues, including market concentration and customer impact.

For Synopsys, acquiring Ansys would potentially enhance its position as a one-stop-shop for engineering and design software, offering integrated solutions to its clients. For Ansys, joining forces with Synopsys could provide access to new markets and customers, leveraging Synopsys's global reach and resources.

In summary, the potential acquisition of Ansys by Synopsys could represent a major shift in the engineering and design software landscape, creating a combined entity with enhanced capabilities and market presence. However, the deal is still under negotiation and subject to various approvals and conditions.

In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)

Use the sliders below to indicate your ratings 1 2 3 4 5 Trust () Prolixity (the use of too many words to express an idea) () Engagement () Substance () Clarity () Alienation () Complexity () Effort (required to read) () Neutrality ()

Human Article - Fina US inflation decelerating in boost to economy

WASHINGTON, Dec 22 (Reuters) - U.S. prices fell in November for the first in more than 3-1/2 years, pushing the annual increase in inflation further below 3%, and boosting financial market expectations of an interest rate cut from the Federal Reserve next March.

The report from the Commerce Department on Friday also showed underlying inflation pressures continuing to subside. Cooling inflation left more income at the disposal of households, helping to underpin consumer spending and the overall economy as the year winds down.

This was yet another data set showcasing the durability of the economic expansion, thanks to a resilient labor market. The economy has defied dire predictions of recession from economists and some business executives going back to late 2022.

"(Fed) Chair (Jerome) Powell couldn't have asked for a better present this year," said Sal Guatieri, a senior economist at BMO Capital Markets in Toronto. "So far at least, the endgame is turning out better than the Fed or nearly anyone could have imagined at the start of the year. While the Fed won't rush into cutting rates, it's likely now just a matter of time."

Inflation, as measured by the personal consumption expenditures (PCE) price index, fell 0.1% last month, the Commerce Department's Bureau of Economic Analysis said. That was the first monthly decline in the PCE price index since April 2020 and followed an unchanged reading in

October.

Food prices edged down 0.1% and energy prices dropped 2.7%. In the 12 months through November, the PCE price index increased 2.6% after rising 2.9% in October. October marked the first time since March 2021 that the annual PCE price index was below 3%.

Economists polled by Reuters had forecast the PCE price index unchanged on the month and rising 2.8% year-on-year.

Excluding the volatile food and energy components, the PCE price index rose 0.1% in November, matching October's gain.

The so-called core PCE price index advanced 3.2% year-on-year, the smallest rise since April 2021, after increasing 3.4% in October. The Fed tracks the PCE price measures for its 2% inflation target.

The government reported on Thursday that core PCE inflation increased at a 2.0% annualized rate in the third quarter. That, combined with November's mild gain, put the six-month core PCE inflation rate at 1.9%.

Monthly inflation readings of 0.2% on a sustainable basis are needed to bring inflation back to the Fed's target, economists say. Financial markets saw a roughly 75% chance of a rate cut at the Fed's March 19-20 policy meeting, according to CME Group's FedWatch Tool.

Subsiding inflation is brightening the mood for many Americans, with a separate report from the University of Michigan on Friday showing consumer sentiment soaring in December, reversing all declines from the previous four months.

President Joe Biden, whose popularity has been hurt by unhappiness over the high cost of living, welcomed the news.

"This reflects the hard work we did together to fix our supply chains and the surge of Americans into the workforce. It's remarkable progress," Biden said in a statement.

Stocks on Wall Street were trading higher. The dollar fell against basket of currencies. U.S. Treasury prices rose.

In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)

Use the sliders below to indicate your ratings 1 2 3 4 5 Trust () Prolixity (the use of too many words to express an idea) () Engagement () Substance () Clarity () Alienation () Complexity () Effort (required to read) () Neutrality ()

Human Article - Poli Supreme Court rejects prosecutor's request, will not rule on Trump immunity yet

WASHINGTON, Dec 22 (Reuters) - The U.S. Supreme Court on Friday declined to immediately decide former President Donald Trump's claim that he cannot be prosecuted for trying to overturn his 2020 election defeat, allowing a lower court to continue reviewing the issue.

The justices, rebuffing an extraordinary request by U.S. Special Counsel Jack Smith, refused to effectively leap-frog a lower appeals court to speed up a final ruling on Trump's claim of criminal immunity ahead of his trial, due to begin in March.

The court issued its decision without explanation in a brief one-page order. No justice publicly dissented. The justices could still decide to consider the issue later.

A federal appeals court in Washington has fast-tracked its consideration of the issue, and scheduled oral arguments for Jan. 9.

In a statement to supporters, Trump labeled the Supreme Court decision as a victory but added "I will still have to fight for my rights in the Appeals Court."

A spokesperson for Smith declined to comment.

Prosecutors have accused Trump, the frontrunner for the 2024 Republican presidential nomination, of attempting to obstruct Congress and defraud the U.S. government through schemes to reverse Democratic President Joe Biden's 2020 election win.

Trump has argued that the case should be dismissed on the grounds that former presidents cannot face criminal charges for conduct related to their official responsibilities.

U.S. District Judge Tanya Chutkan rejected that claim on Dec. 1, prompting Trump's appeal to the U.S. Court of Appeals for the District of Columbia Circuit. Trump's appeal suspended his

trial, currently set for March.

In a bid to avoid delaying Trump's trial, Smith on Dec. 11 urged the Supreme Court to undertake an expedited ruling - even as the D.C. Circuit court races to rule on the issue.

If Trump is reelected to the White House on Nov. 5, he could seek to pardon himself of any federal crimes.

Trump's lawyers opposed Smith's request, telling the justices in a court filing the special counsel's bid to bypass the normal appellate process amounted to a "rush to decide the issues with reckless abandon."

Three of the nine Supreme Court justices were appointed by Trump, who cemented a 6-3 conservative majority.

In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)

Use the sliders below to indicate your ratings 1 2 3 4 5 Trust () Prolivity (the use of too many words to express an idea) () Engagement () Substance () Clarity () Alienation () Complexity () Effort (required to read) () Neutrality ()

Human Article - Tech Exclusive: Synopsys seeks to acquire engineering software company Ansys-sources

Dec 22 (Reuters) - (This Dec. 22 story has been corrected to fix the reference to Ansys, not Synopsys, in paragraph 8)

Synopsys Inc (SNPS.O), a maker of software used in chip design, has submitted an offer to acquire Ansys Inc (ANSS.O), an engineering software vendor with a market value of \$30 billion, people familiar with the matter said on Friday.

The deal negotiations come as Synopsys co-founder Aart de Geus prepares to transition to an executive chairman role and hand over the chief executive reins to chief operating officer Sassine Ghazi on Jan. 1.

Synopsys' pursuit of a transformative acquisition amid a leadership change underscores the commercial appeal of Ansys' software, which is used widely in design, including in tennis to design rackets for players like Novak Djokovic.

The offers that Ansys has attracted value it at well over \$400 per share, and a deal could be announced in the coming weeks if the talks do not fall apart, the sources said.

The sources requested anonymity as the discussions are confidential. Ansys and Synopsys declined to comment.

Bloomberg News earlier reported on Thursday Ansys was exploring a sale, without naming any suitor.

Ansys shares ended trading in New York on Friday up 18% at \$357.98. Synopsys shares dropped 6% to \$524.46.

Ansys is a maker of simulation software, which is typically used to forecast how it can be applied across different industries in the real world.

Synopsys, which has a market value of \$85 billion, supplies software that engineers use to design and test chip circuits.

In the following question, please rate how well the article satisfies the following criteria, from "meets very poorly" (1) to "meets very well" (5)

Use the sliders below to indicate your ratings 1 2 3 4 5 Trust () Prolixity (the use of too many words to express an idea) () Engagement () Substance () Clarity () Alienation () Complexity () Effort (required to read) () Neutrality ()

Here is your ID: `{e://Field/Random%20ID}`

Copy this value to paste into MTurk. When you have copied the ID, please click the next button to submit your survey.

D) Unfiltered List of Descriptors (Phase 1)

- 1. Amazing**
- 2. Balanced**
- 3. Believable**
- 4. Boring**
- 5. Complete**

6. Complex

7. Concise

8. Confrontational - appropriateness (reaction to subject matter)

9. Decisive - redundancy (conclusive)

10. Definitive - redundancy (conclusive)

11. Detailed

12. Digital - appropriateness (reaction to subject matter - digital rights act article)

13. Digital - appropriateness (digital rights act article)

14. Disappointing - appropriateness (reaction to subject matter)

15. Educational

16. Fabulous - redundancy (amazing)

17. Factual

18. Factual - duplicate

19. Flowing - **changed to Coherent**

20. Great - redundancy (amazing)

21. Greedy - appropriateness (reaction to subject matter)

22. Honest

23. Inconclusive (Conclusive)

24. Informative

25. Informative - dup

26. Informative - dup

27. Informative - dup

28. Informative - dup

29. Interesting

30. Interesting - dup

31. Lengthy - redundancy (long)

32. Limiting - appropriateness (reaction to subject matter)

33. Long

34. Long - duplicate

35. Long - duplicate

36. Moderate

37. Moral

38. Neutral

39. New - redundancy (timely)

40. Possible - appropriateness (subj matter)

41. Precise

42. Rote

43. Sad - appropriateness (subj matter)

44. Technical

45. Thorough

46. Thought-provoking

47. Thought-Provoking - duplicate

48. Threatening - appropriateness (subj matter)

49. Timely

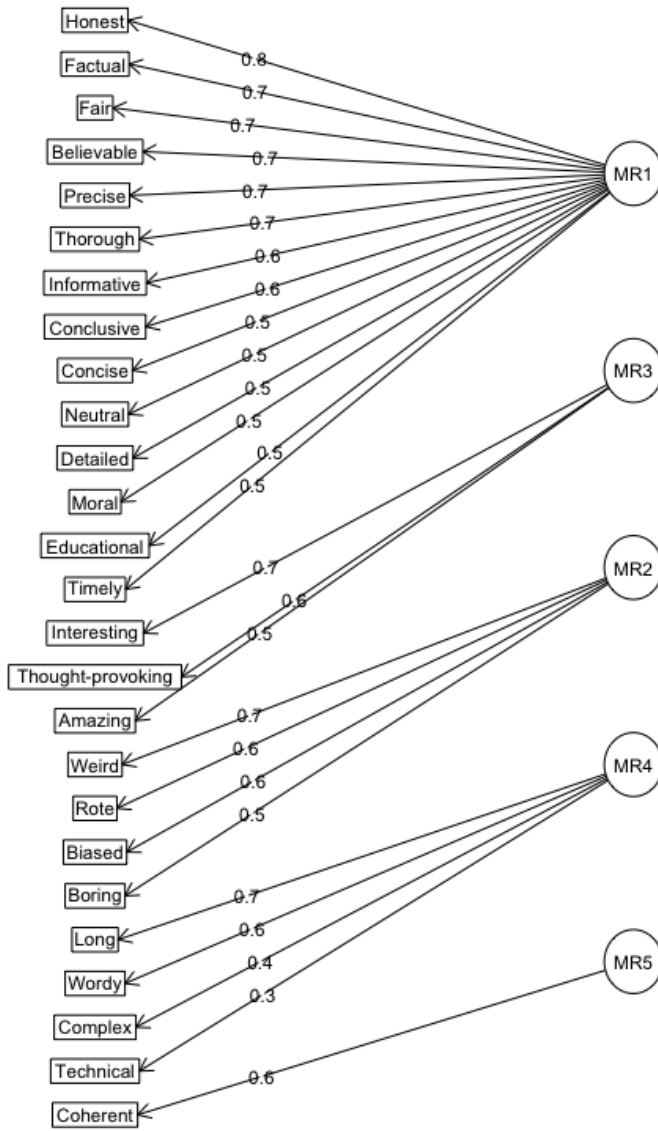
50. Unbelievable - redundancy (believable)

51. Unbiased (Biased)

- 52. Unbiased - duplicate
- 53. Uncompromising - appropriateness (subj matter)
- 54. Unfair - redundancy
- 55. Uninteresting - redundancy
- 56. Verbose - redundancy
- 57. Weird**
- 58. Wordy**
- 59. Wrong - redundancy (factual)

E) EFA Results Diagram - Varimax

EFA Results - Varimax Rotation



F) EFA Factors - Varimax

Factor 1 - Quality

Honest

Factual

Fair

Believable

Precise

Thorough

Informative

Conclusive

Concise

Neutral

Detailed

Moral

Educational

Timely

Factor 2 - Engagement

Interesting

Thought-provoking

Amazing

Factor 3 - Alienation

Weird

Rote

Biased

Boring

Factor 4 - **Effort**

Long

Wordy

Complex

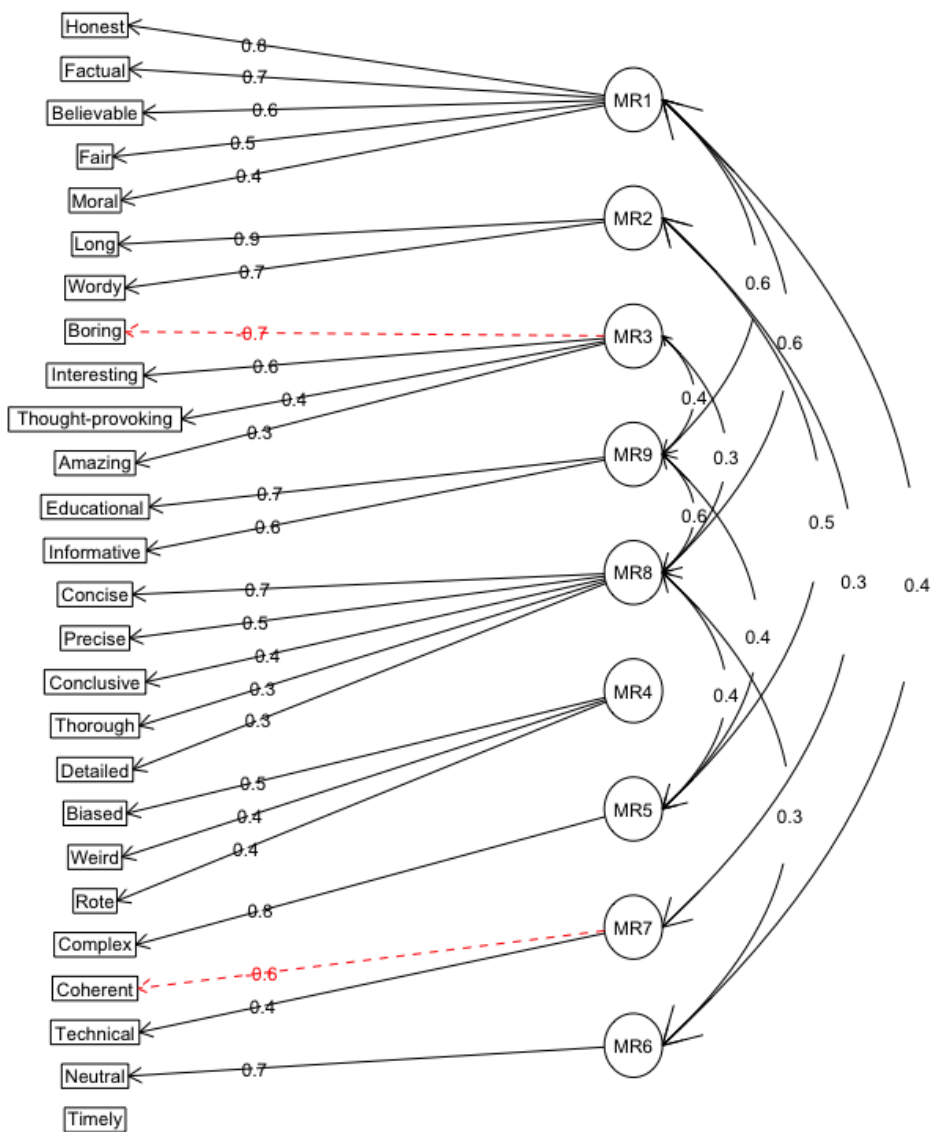
Technical

Factor 5 - **Coherence**

Coherent

G) EFA Results Diagram - Oblimin

EFA Results - Oblimin Rotation



H) EFA Factors - Oblimin

Factor 1 (MR1) - Trust

Honest

Factual

Believable

Fair

Moral

Factor 2 (MR2) - Prolixity

Long

Wordy

Factor 3 (MR3) - Engagement

Boring

Interesting

Thought-provoking

Amazing

Factor 4 (MR9) - Substance

Educational

Informative

Factor 5 (MR8) - Clarity

Concise

Precise

Conclusive

Thorough

Detailed

Factor 6 (MR4) - Alienation

Biased

Weird

Rote

Factor 7 (MR5) - Complexity

Complex

Factor 8 (MR7) - Effort

Coherent

Technical

Factor 9 (MR6) - Neutrality

Neutral

No factor: Timely

References

- Adeoye-Olatunde, Omolola A., and Nicole L. Olenik. "Research and scholarly methods: Semi-structured interviews." *Journal of the American College of Clinical Pharmacy* 4, no. 10 (2021): 1358-1367.
- Akst, Daniel. "Automation anxiety." *The Wilson Quarterly* 37, no. 3 (2013): 65.
- Amoore, Louise. *Cloud ethics: Algorithms and the attributes of ourselves and others*. Duke University Press, 2020.
- AP NEWS. "Alibaba: Fiscal Q4 Earnings Snapshot," May 13, 2021.
<https://apnews.com/article/earnings-business-52ba246a097ab4f79d372c39812f37fa>.
- Autor, David. "Why are there still so many jobs? The history and future of workplace automation." *Journal of economic perspectives* 29, no. 3 (2015): 3-30.
- Barr, Jeremy. "Mehdi Hasan saw a market for a new kind of media company. So far, so good." *Washington Post*. Sep 4 2024.
<https://www.washingtonpost.com/style/media/2024/09/04/mehdi-hasan-zeteo-substack/>
- Bates, Marcia J. "Information behavior." *Encyclopedia of library and information sciences* 3 (2010): 2381-2391.
- Bauder, David. "Sports Illustrated found publishing AI generated stories, photos, and authors". *PBS.com*. Nov 29, 2023.
<https://www.pbs.org/newshour/economy/sports-illustrated-found-publishing-ai-generate-stories-photos-and-authors>
- Beckett, Charlie. "Preparing for the coming wave of generative AI in journalism." *London School of Economics Impact Blog*. 18 Sep 2023.

<https://blogs.lse.ac.uk/impactofsocialsciences/2023/09/18/preparing-for-the-coming-wave-of-generative-ai-in-journalism/>

Benson, Rodney. "Field theory in comparative context: A new paradigm for media studies."

Theory and society 28, no. 3 (1999): 463-498.

Berry, David. "The computational turn: Thinking about the digital humanities." *Culture machine* 12 (2011).

Boczkowski, Pablo J. "Materiality and mimicry in the journalism field." In *The changing faces of journalism*, pp. 66-77. Routledge, 2009.

Bourdieu, Pierre. *The field of cultural production: Essays on art and literature*. Columbia University Press, 1993.

Boyd, Danah, and Kate Crawford. "Critical questions for big data: Provocations for a cultural, technological, and scholarly phenomenon." *Information, communication & society* 15, no. 5 (2012): 662-679.

Brown, John Seely, and Paul Duguid. *The social life of information: Updated, with a new preface*. Harvard Business Review Press, 2017.

Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan et al. "Language models are few-shot learners." *arXiv preprint arXiv:2005.14165* (2020).

Buckland, Michael. (1997). "What is a Document?" Retrieved from

<http://people.ischool.berkeley.edu/~buckland/whatdoc.html>

Buckland, Michael Keeble. *Information and Society*. The MIT Press, 2017.

Burrell, Jenna. "How the machine 'thinks': Understanding opacity in machine learning algorithms." *Big Data & Society* 3, no. 1 (2016): 2053951715622512.

- Bussolini, Jeffrey. "What is a Dispositive?." *Foucault studies* (2010): 85-107.
- Cahn, Albert Fox. "A Human Wrote This Article. You Shouldn't Be Scared of GPT-3" *The Guardian*. Sep 12 2020.
<http://www.theguardian.com/commentisfree/2020/sep/12/human-wrote-this-article-gpt-3>.
- Carlson, Matt. "The robotic reporter: Automated journalism and the redefinition of labor, compositional forms, and journalistic authority." *Digital journalism* 3, no. 3 (2015): 416-431.
- Chaffee, Steven H., and Joan Schleuder. "Measurement and effects of attention to media news." *Human communication research* 13, no. 1 (1986): 76-107.
- Chaney, Allison JB, Brandon M. Stewart, and Barbara E. Engelhardt. "How algorithmic confounding in recommendation systems increases homogeneity and decreases utility." In *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224-232. 2018.
- Chen, Yimin, Niall J. Conroy, and Victoria L. Rubin. "Misleading online content: recognizing clickbait as" false news"." In *Proceedings of the 2015 ACM on workshop on multimodal deception detection*, pp. 15-19. 2015.
- Cho, Kyunghyun, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. "Learning phrase representations using RNN encoder-decoder for statistical machine translation." *arXiv preprint arXiv:1406.1078* (2014).
- Chuang, Tyng-Ruey, Rebecca C. Fan, Ming-Syuan Ho, and Kalpana Tyagi. "Openness." *Internet Policy Review* 11, no. 1 (2022): 1-9.
- Coglianesi, Cary, and David Lehr. "Transparency and algorithmic governance." *Admin. L. Rev.* 71 (2019): 1.

Coldewey, David. "OpenAI shifts from nonprofit to 'capped-profit' to attract capital"

TechCrunch. Mar 11 2019.

<https://techcrunch.com/2019/03/11/openai-shifts-from-nonprofit-to-capped-profit-to-attract-capital/>

Cooper, Mark N. "The case against media consolidation." *New York: Donald McGannon Communication Research Center* (2007).

Dale, Robert. "GPT-3: What's it good for?" *Toward Data Science*. 8 Dec 2020.

<https://towardsdatascience.com/gpt-3-whats-it-good-for-156a445cefc8>

Danzon-Chambaud, Samuel. "A systematic review of automated journalism scholarship: guidelines and suggestions for future research." *Open Research Europe* 1, no. 4 (2021): 4.

Danzon-Chambaud, Samuel. "Automated news in practice: changing the journalistic doxa during COVID-19, at the BBC and across media organisations." PhD diss., Dublin City University, 2023.

Daume, Harold Charles, and Daniel Marcu. *Practical structured learning techniques for natural language processing*. University of Southern California, 2006.

de Haan, Yael, Eric van den Berg, Nele Goutier, Sanne Kruikemeier, and Sophie Lecheler.

"Invisible friend or foe? How journalists use and perceive algorithmic-driven tools in their research process." *Digital Journalism* 10, no. 10 (2022): 1775-1793.

Diakopoulos, Nicholas. *Automating the news: How algorithms are rewriting the media*. Harvard University Press, 2019.

Diakopoulos, Nick. "What Could ChatGPT Do for News Production?" *Medium*. 15 Feb 2023.

<https://generative-ai-newsroom.com/what-could-chatgpt-do-for-news-production-2b2a3b7047d9>

- Dörr, Konstantin Nicholas. "Mapping the field of algorithmic journalism." *Digital journalism* (2015).
- Elkins, Katherine, and Jon Chun. "Can GPT-3 pass a writer's Turing test?." *Journal of Cultural Analytics* 5, no. 2 (2020).
- Elon Musk, Twitter post, Sep 24, 2020, 1:50AM PST,
<https://twitter.com/elonmusk/status/1309052632850468864>
- Fish, Stanley E. "Interpreting the "variorum"." *Critical inquiry* 2, no. 3 (1976): 465-485.
- Floridi, Luciano. "AI as agency without intelligence: on ChatGPT, large language models, and other generative models." *Philosophy & technology* 36, no. 1 (2023): 15.
- Floridi, Luciano, and Massimo Chiriatti. "GPT-3: Its nature, scope, limits, and consequences." *Minds and Machines* 30 (2020): 681-694.
- Friedman, Ellen G., and Miriam Fuchs. "Contexts and Continuities: An Introduction to Women's Experimental Fiction in English." *Breaking the Sequence: Women's Experimental Fiction* (1989): 3-51.
- Furner, Jonathan. "Information studies without information." *Library Trends*, 52, no. 3 (2004): 427-446.
- Galtung, Johan, and Mari Holmboe Ruge. "The structure of foreign news: The presentation of the Congo, Cuba and Cyprus crises in four Norwegian newspapers." *Journal of peace research* 2, no. 1 (1965): 64-90.
- Ghioni, Riccardo, Mariarosaria Taddeo, and Luciano Floridi. "Open source intelligence and AI: a systematic review of the GELSI literature." *AI & society* 39, no. 4 (2024): 1827-1842.
- Goertzel, Ben. "Artificial general intelligence: concept, state of the art, and future prospects." *Journal of Artificial General Intelligence* 5, no. 1 (2014): 1-48.

Goodwin, Jazmin. "Elon Musk criticizes OpenAI exclusively licensing GPT-3 to Microsoft".

CNN Business. Sep 27 2020.

<https://edition.cnn.com/2020/09/27/tech/elon-musk-tesla-bill-gates-microsoft-open-ai/index.html>

Gorman, Gary Eugene, Peter Robert Clayton, Sydney J. Shep, and Adela Clayton. *Qualitative research for the information professional: A practical handbook*. Facet Publishing, 2005.

Graefe, Andreas, and Nina Bohlken. "Automated journalism: A meta-analysis of readers' perceptions of human-written in comparison to automated news." *Media and Communication* 8, no. 3 (2020): 50-59.

Graefe, Andreas, and Nina Bohlken. "Automated Journalism: A Meta-Analysis of Readers' Perceptions of Human-Written in Comparison to Automated News." *Media and Communication* 8, no. 3 (2020): 50-59.

Graefe, Andreas. "Guide to automated journalism." *Columbia Academic Commons* (2016).

<https://academiccommons.columbia.edu/doi/10.7916/D80G3XDJ>

Green, Ben. "The false promise of risk assessments: epistemic reform and the limits of fairness." In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 594-606. 2020.

Guillory, John. "The memo and modernity." *Critical inquiry* 31, no. 1 (2004): 108-132.

Harmon, Joanie. "Ramesh Srinivasan on ChatGPT" *UCLA School of Education and Information Studies*. Feb 21, 2023. <https://seis.ucla.edu/news/ramesh-srinivasan-on-chatgpt/>

Hashmi, Sahar, MD, PhD. "2023 Is the Era of Generative AI" *Entrepreneur*. Mar 10, 2023.

<https://www.entrepreneur.com/science-technology/2023-is-the-era-of-ai-so-whats-in-it-for-entrepreneurs/444633>

Heaven, Will. "Generative AI is changing everything. But what's left when the hype is gone?."

MIT Technology Review. Dec 16, 2022.

<https://www.technologyreview.com/2022/12/16/1065005/generative-ai-revolution-art/>

Henestrosa, Angelica Lermann, Hannah Greving, and Joachim Kimmerle. "Automated journalism: The effects of AI authorship and evaluative information on the perception of a science journalism article." *Computers in Human Behavior* 138 (2023): 107445.

Hesmondhalgh, David. "Bourdieu, the media and cultural production." *Media, culture & society* 28, no. 2 (2006): 211-231.

Hobart, Michael E., and Zachary S. Schiffman. *Information ages: Literacy, numeracy, and the computer revolution*. JHU Press, 2000.

Howard, Jeremy, and Sebastian Ruder. "Universal language model fine-tuning for text classification." arXiv preprint arXiv:1801.06146 (2018).

Huang, J. Sonia, and Don Heider. "Media convergence: A case study of a cable news station." *The International Journal on Media Management* 9, no. 3 (2007): 105-115.

Ihalainen, Jani. "Computer creativity: artificial intelligence and copyright." *Journal of Intellectual Property Law & Practice* (2018).

Jeon, Doh-Shin, and Nikrooz Nasr. "News aggregators and competition among newspapers on the internet." *American Economic Journal: Microeconomics* 8, no. 4 (2016): 91-114.

Joye, S., Heinrich, A., & Wöhlert, R. (2016). 50 years of Galtung and Ruge: Reflections on their model of news values and its relevance for the study of journalism and communication today. *CM. Communication and Media*, 11(36), 5-28.

Karbal, Ian. "Artificial Intelligence that can write stories and crunch data is spreading in newsrooms. That's a good thing for journalists." *Gateway Journalism Review* 47, no. 352

(2019): 24-27.

Katz, Daniel Martin, Michael James Bommarito, Shang Gao, and Pablo Arredondo. "Gpt-4 passes the bar exam." *Philosophical Transactions of the Royal Society A* 382, no. 2270 (2024): 20230254.

Katzenbach, Christian, and Lena Ulbricht. "Algorithmic governance." *Internet Policy Review* 8, no. 4 (2019): 1-18.

Knibbs, Kate. "Confessions of an AI Clickbait Kingpin". *WIRED Magazine*. Feb 7 2024.

<https://www.wired.com/story/confessions-of-an-ai-clickbait-kingpin/>

Knott, Eleanor, Aliya Hamid Rao, Kate Summers, and Chana Teeger. "Interviews in the social sciences." *Nature Reviews Methods Primers* 2, no. 1 (2022):

73.<https://doi.org/10.1038/s43586-022-00150-6>

Koopman, Colin. *How we became our data: A genealogy of the informational person*. University of Chicago Press, 2019.

Kotsiantis, Sotiris B., I. Zaharakis, and P. Pintelas. "Supervised machine learning: A review of classification techniques." *Emerging artificial intelligence applications in computer engineering* 160, no. 1 (2007): 3-24.

Koubaa, Anis. "GPT-4 vs. GPT-3.5: A concise showdown." (2023).

Kruzel, John. "Supreme Court rejects prosecutor's request, will not rule on Trump immunity yet" *Reuters*. Dec 22 2023.

<https://www.reuters.com/legal/us-supreme-court-wont-rule-trumps-immunity-claim-2020-election-case-now-2023-12-22/>

Leiser, M. R. "Bias, journalistic endeavours, and the risks of artificial intelligence." In *Artificial Intelligence and the Media*, pp. 8-32. Edward Elgar Publishing, 2022.

Leppänen, Leo, Myriam Munezero, Mark Granroth-Wilding, and Hannu Toivonen. "Data-driven news generation for automated journalism." In *Proceedings of the 10th international conference on natural language generation*, pp. 188-197. 2017.

<https://doi.org/10.18653/v1/w17-3528>

Li, Chuan. "OpenAI's GPT-3 Language Model: A Technical Overview" *LambdaLabs*. Jun 3 2020. <https://lambdalabs.com/blog/demystifying-gpt-3>

Linden, Carl-Gustav. "Decades of Automation in the Newsroom: Why are there still so many jobs in journalism?." *Digital journalism* 5, no. 2 (2017): 123-140.

Maack, Mary Niles. "The lady and the antelope: Suzanne Briet's contribution to the French documentation movement." (2004).

Marr, Bernard. "A Short History Of ChatGPT: How We Got To Where We Are Today" *Forbes*. May 19 2023.

<https://www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/>

Martinez, Gina. "Potential Category 4 atmospheric river to hit Northern California" *CBS News*. Dec 27 2022.

<https://www.cbsnews.com/news/northern-california-pacific-northwest-atmospheric-river-today-2022-12-27/>

Martinez, Maria. "Volkswagen hit by IT outage, brand vehicle production in Germany halted" *Reuters*. Sep 27 2023.

<https://www.reuters.com/business/autos-transportation/it-disruption-paralyses-volkswagens-central-infrastructure-2023-09-27/>.

Marvin, Ggaliwango, Nakayiza Hellen, Daudi Jjingo, and Joyce Nakatumba-Nabende. "Prompt

- engineering in large language models." In *International conference on data intelligence and cognitive informatics*, pp. 387-402. Singapore: Springer Nature Singapore, 2023.
https://doi.org/10.1007/978-981-99-7962-2_30
- Maxwell, Joseph A. *Qualitative research design: An interactive approach*. Sage publications, 2012.
- McChesney, Robert W., and John Nichols. *The death and life of American journalism: The media revolution that will begin the world again*. Bold Type Books, 2011.
- McGrath, Richard. "Journalling and memoing: Reflexive qualitative research tools." In *Handbook of qualitative research methodologies in workplace contexts*, pp. 245-262. Edward Elgar Publishing, 2021.
- McGurl, Mark. "The Program Era: Pluralisms of Postwar American Fiction." *Critical Inquiry* 32, no. 1 (2005): 102-129.
- Mencher, Melvin, and Wendy P. Shilton. *News reporting and writing*. Madison, WI: Brown & Benchmark Publishers, 1997.
- Metz, Cade. "OpenAI Completes Deal That Values the Company at \$80 billion" *New York Times*. Feb 16 2024.
<https://www.nytimes.com/2024/02/16/technology/openai-artificial-intelligence-deal-valuation.html>
- Miller, Ross. "AP's 'robot journalists' are writing their own stories now." *The Verge*. Jan 29 2015.
<https://www.theverge.com/2015/1/29/7939067/ap-journalism-automation-robots-financial-reporting>
- Montal, Tal, and Zvi Reich. "I, robot. You, journalist. Who is the author? Authorship, bylines and full disclosure in automated journalism." *Digital journalism* 5, no. 7 (2017): 829-849.

Mor, Sahar. "Can GPT-3 Build a GPT-3 App?" *Medium*. Sep 7 2020.

<https://medium.com/swlh/can-gpt-3-build-a-gpt-3-app-dc4d17a5b351>

Müller, Berndt, Joachim Reinhardt, and Michael T. Strickland. *Neural networks: an introduction*. Springer Science & Business Media, 2012.

Mutikani, Lucia. "US inflation decelerating in boost to economy" *Reuters*. Dec 22 2023.

<https://www.reuters.com/markets/us/us-annual-inflation-slows-further-below-3-novembe-2023-12-22/>.

Napoli, Philip M. "Automated media: An institutional theory perspective on algorithmic media production and consumption." *Communication Theory* 24, no. 3 (2014): 340-360.

Naumer, Charles, Karen Fisher, and Brenda Dervin. "Sense-Making: a methodological perspective." In *Sensemaking Workshop, CHI*, vol. 8. 2008.

Nellis, Stephen. "Epic Games asks US Supreme Court to review Apple antitrust case" *Reuters*. Sep 27 2023.

<https://www.reuters.com/legal/epic-games-asks-us-supreme-court-review-apple-antitrust-case-2023-09-27/>

Nielsen, Jakob. "AI: First New UI Paradigm in 60 Years" *Norman-Nielsen Group News*. Jun 18 2023. <https://www.nngroup.com/articles/ai-paradigm/>

Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism*. NYU Press, 2018.

Nudelman, Greg. *Designing search: UX strategies for ecommerce success*. John Wiley & Sons, 2011.

Nunnally, Jum C. "An overview of psychological measurement." *Clinical diagnosis of mental disorders: A handbook* (1978): 97-146.

- Oliffe, John L., Mary T. Kelly, Gabriela Gonzalez Montaner, and Wellam F. Yu Ko. "Zoom interviews: Benefits and concessions." *International journal of qualitative methods* 20 (2021): 16094069211053522.
- Panagia, Davide. "ChatGPT and the Technical Ontologies of Critical Political Science." *New Political Science* 45, no. 3 (2023): 555-558.
- Panagia, Davide. "On the possibilities of a political theory of algorithms." *Political Theory* 49, no. 1 (2021): 109-133.
- Panagia, Davide. "The Algorithm Dispositif: Risk and Automation in the Age of# datapolitik." In *The Routledge Companion to Media and Risk*, pp. 118-129. Routledge, 2020.
- Pariser, Eli. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.
- Pasquale, Frank. *The black box society*. Harvard University Press, 2015, 2
- Peng, Gang, and Rahul Bhaskar. "Artificial intelligence and machine learning for job automation: A review and integration." *Journal of Database Management (JDM)* 34, no. 1 (2023): 1-12.
- Pickard, Alison Jane. *Research methods in information*. Facet publishing, 2013.
- Porter, Michael E., and Mark R. Kramer. "Advantage." *Creating and Sustaining Superior Performance*, Simons (1985): 56-68.
- Porter, Michael E. "The value chain and competitive advantage." *Understanding business processes* 2 (2001): 50-66.
- Radford, Alec, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. "Improving language understanding by generative pre-training." (2018).
- Raymond, Nate and Blake Brittain. "Internet Archive's digital book lending violates copyrights,

US judge rules” *Reuters*. Mar 27 2023.

<https://www.reuters.com/legal/us-judge-rules-internet-archives-digital-book-lending-violates-copyrights-2023-03-25/>

Reid, Tim and Joseph Ax. “Time running out for Trump’s Republican rivals as they face off in second debate” *Reuters*. Sep 28 2023.

<https://www.reuters.com/world/us/time-running-out-trumps-republican-2024-rivals-they-face-off-second-debate-2023-09-27/>

Rogers, Don. “Journalism Matters: The canary in the coal mine” *The Park Record*. Dec 5 2023.

<https://www.parkrecord.com/2023/12/05/journalism-matters-the-canary-in-the-coal-mine/>

Rosoł, Maciej, Jakub S. Gąsior, Jonasz Łaba, Kacper Korzeniewski, and Marcel Młyńczak.

"Evaluation of the performance of GPT-3.5 and GPT-4 on the Polish Medical Final Examination." *Scientific Reports* 13, no. 1 (2023): 20512.

Sadler, Georgia Robins, Hau-Chen Lee, Rod Seung-Hwan Lim, and Judith Fullerton.

"Recruitment of hard-to-reach population subgroups via adaptations of the snowball sampling strategy." *Nursing & health sciences* 12, no. 3 (2010): 369-374.

Sathya, Ramadass, and Annamma Abraham. "Comparison of supervised and unsupervised

learning algorithms for pattern classification." *International Journal of Advanced*

Research in Artificial Intelligence 2, no. 2 (2013): 34-38.

Seabrook, John. “The Next Word.” *The New Yorker*. Oct 14 2019.

<https://www.newyorker.com/magazine/2019/10/14/can-a-machine-learn-to-write-for-the-new-yorker>

Sen, Anirban and Vinn, Milana. “Exclusive: Synopsys seeks to acquire engineering software company Ansys-sources” *Reuters*. Dec 24 2023.

<https://www.reuters.com/markets/deals/ansys-explores-sale-amid-takeover-interest-sources-2023-12-22/>

Shannon, Claude E. "A mathematical theory of communication." *The Bell system technical journal* 27, no. 3 (1948): 379-423.

Simondon, Gilbert. "Du mode d'existence des objets techniques." (1989).

Spohr, Dominic. "Fake news and ideological polarization: Filter bubbles and selective exposure on social media." *Business Information Review* 34, no. 3 (2017): 150-160.

Sundar, S. S., & Kim, J. (2019, May). Machine heuristic: When we trust computers more than humans with our personal information. In *Proceedings of the 2019 CHI Conference on human factors in computing systems* (pp. 1-9).

Sundar, S. Shyam, Anne Oeldorf-Hirsch, and Amulya Garga. "A cognitive-heuristics approach to understanding presence in virtual environments." In *PRESENCE 2008: Proceedings of the 11th Annual International Workshop on Presence*, pp. 219-228. Padova, Italy: CLEUP Cooperativa Libreria Universitaria Padova, 2008.

Sundar, S. Shyam. "Exploring receivers' criteria for perception of print and online news." *Journalism & Mass Communication Quarterly* 76, no. 2 (1999): 373-386.

Tausk, Victor. "On the origin of the "influencing machine" in schizophrenia." *The Psychoanalytic Quarterly* 2, no. 3-4 (1933): 519-556.

The Economist. "How to respond to climate change, if you are an algorithm". *The Economist*. Oct 1 2019.

<https://www.economist.com/open-future/2019/10/01/how-to-respond-to-climate-change-if-you-are-an-algorithm>

Thurman, Neil, Konstantin Dörr, and Jessica Kunert. "When reporters get hands-on with

- robo-writing: Professionals consider automated journalism's capabilities and consequences." *Digital journalism* 5, no. 10 (2017): 1240-1259.
- Tucker, Joshua A., Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. "Social media, political polarization, and political disinformation: A review of the scientific literature." *Political polarization, and political disinformation: a review of the scientific literature (March 19, 2018)* (2018).
- Van Dalen, Arjen. "The algorithms behind the headlines: How machine-written news redefines the core skills of human journalists." *Journalism practice* 6, no. 5-6 (2012): 648-658.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. "Attention is all you need." *arXiv preprint arXiv:1706.03762* (2017).
- Verma, Nitin, Kenneth R. Fleischmann, and Kolina S. Koltai. "Human values and trust in scientific journals, the mainstream media and fake news." *Proceedings of the Association for Information Science and Technology* 54, no. 1 (2017): 426-435.
- Vincent, James. "Microsoft invests \$1 billion in OpenAI to pursue holy grail of artificial intelligence". *The Verge*. Jul 22 2019.
<https://www.theverge.com/2019/7/22/20703578/microsoft-openai-investment-partnership-1-billion-azure-artificial-general-intelligence-agi>
- Wachter, Sandra. "Limitations and loopholes in the EU AI Act and AI Liability Directives: what this means for the European Union, the United States, and beyond." *Yale Journal of Law and Technology* 26, no. 3 (2024).
- Walter, Yoshija. "Artificial influencers and the dead internet theory." *AI & SOCIETY* (2024): 1-2.
- Wang, Sai, and Guanxiong Huang. "The impact of machine authorship on news audience

- perceptions: a meta-analysis of experimental studies." *Communication Research* (2024): 00936502241229794.
- Weise, Karen. "Inside the A.I. Arms Race That Changed Silicon Valley Forever" *New York Times*. Dec 5 2023.
<https://www.nytimes.com/2023/12/05/technology/ai-chatgpt-google-meta.html>
- Wiggers, Kyle. "The New York Times wants OpenAI and Microsoft to pay for training data" *TechCrunch*. 27 Dec 2023.
<https://techcrunch.com/2023/12/27/the-new-york-times-wants-openai-and-microsoft-to-pay-for-training-data/>
- Will, Barbara. "Gertrude Stein, automatic writing and the mechanics of genius." In *Forum for modern language studies*, vol. 37, no. 2, pp. 169-175. Oxford University Press, 2001.
- Wu, Shangyuan, Edson C. Tandoc, and Charles T. Salmon. "A Field Analysis of Journalism in the Automation Age: Understanding Journalistic Transformations and Struggles Through Structure and Agency." *Digital Journalism* 7, no. 4 (2019): 428-446.
- Yang, Zhilin, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. "Xlnet: Generalized autoregressive pretraining for language understanding." *arXiv preprint arXiv:1906.08237* (2019).
- York, Jessica A. "What's that giant orb in Live Oak?" *Santa Cruz Sentinel*. Sep 10 2022.
<https://www.santacruzsentinel.com/2022/09/10/whats-that-giant-orb/>
- Zuboff, Shoshana. "Automate/informate: The two faces of intelligent technology." *Organizational dynamics* 14, no. 2 (1985): 5-18.