Spring 5-1-2019

AI in Journalism: Creating an Ethical Framework

Haley Kim

AI in Journalism: Creating an Ethical Framework

A Thesis Project Submitted in Partial Fulfillment of the
Requirements of the Renée Crown University Honors Program at
Syracuse University

Haley Kim

Candidate for Bachelor of Science Degree (see note)
and Renée Crown University Honors
Spring 2019

Honors Thesis Project in Magazine Journalism

Thesis Project Advisor: _____
Aileen Gallagher, Associate Professor of Magazine, News and Digital Journalism

Thesis Project Reader: _____
Sean Branagan, Director of the Center for Digital Media Entrepreneurship

Honors Director: _____
Danielle Taana Smith, Director

Abstract

This thesis is an examination of the ethical use of artificial intelligence in journalism. Artificial intelligence is currently being used in all steps of the news production process: story discovery, story production and story distribution. Newsrooms utilize machine learning to analyze massive quantities of data and discover patterns that humans would normally never be able to pick up. Additionally, journalists also create templates so computers can write stories that are data-based, such as earning reports and game (sports) stories, and free them up to be able to work on other projects. Newsrooms can also use AI to personalize story recommendations to readers. While there is great potential for machine learning and AI in journalism, it is also an emerging technology that creates new ethical challenges for newsrooms. Interviews conducted with 12 people working in journalism, technology and law focus on issues of bias, transparency, legislation and attribution for algorithms, among others. Based off this research, an ethical framework was built for newsrooms to follow as they implement this technology.

Executive Summary

This report explores the use of artificial intelligence in journalism and the ethical decision-making process that should come with implementing this technology in U.S. newsrooms. These national, mostly large media organizations use artificial intelligence in three key areas of the news production process: story discovery, story generation and story distribution. Newsrooms that focus on speed, like the *Associated Press* and *Bloomberg*, employ AI in the story production step with automated content generation. Smaller newsrooms have focused more on using AI as a story discovery and reporting tool. Interviews conducted with 12 people who work in media, tech and law focused on the way their organizations incorporated ethical decision-making as they implemented their technology.

Some key findings:

- Many of the ethical issues facing artificial intelligence in journalism are issues that journalists already face, just in a different context. One example is the importance of having clean, reliable data, which is also imperative to data journalism. Another is the need for transparency and disclosure policies for automated content generated pieces, which media organizations already need to do for projects like investigative reporting.
- There are ethics codes for journalism and there are ethics codes for AI, but there are no widely known ethics codes for artificial intelligence use in journalism. Individual media organizations have certainly discussed ethical use of this technology, but few, if any at all, have added these concerns to their codes of ethics.
- Other ethical situations, such as attribution and transparency, are specific to AI. There are not yet industry-wide standards for properly crediting or attributing content "written" via algorithm.

- When it comes to legislation, most interviewed stand on the side of less regulation for reasons such as legislators' lack of knowledge about this technology and the need for organizations to implement their own transparency measures first. Regulations, however, are not the same as ethics, and it's important that media organizations have an ethical code to follow regardless of legislation.

This report proposes an ethical framework that media organizations can follow as they integrate artificial intelligence in their newsrooms. These guidelines focus on steps to consistently evaluate quality of data and algorithms, analyze potential bias in models and act transparently with AI-based endeavors. This framework was developed in conjunction with universal journalism values such as accuracy, fairness, respect and responsibility to the public (Society of Professional Journalists, n.d.).

Table of Contents

Acknowledgements

This thesis was a year-long project that I could not have done without the help of many other people. Many thanks to my thesis advisor, Professor Aileen Gallagher, who answered all my big, little, important and non-important questions, helped me secure funding, edited multiple times and patiently guided me through this project. Thanks to Professor Sean Branagan, my thesis reader, for guiding me through the whirlwind that is South by Southwest and giving critical feedback, especially related to the nature of the technology and the media industry. Thank you to the Renée Crown University Honors Program for allowing me to explore this topic, answering all my logistical questions and funding this project — I am unbelievably grateful for the opportunity to travel. Thank you to everyone I spoke to, whether I interviewed you about your work, your organization or just asked for your feedback — you all took time out of your busy schedules to talk with me. And finally, thank you to my parents, family and friends, who provided emotional support for me as I endeavored on this project for the last year.

Introduction

The Future Institute Today's 2019 Trend Report declared artificial intelligence not just a tech trend, but the third era of computing (Webb, Giralt, Levkowitz, Palatucci, & Perez, 2019). "Artificial intelligence" is a buzzword in media, but what it truly means is more nuanced than just a machine imitating humans.

Artificial intelligence is a system that makes autonomous decisions that appear like acts of human intelligence (Webb, et. al, 2019). There are two types of artificial intelligence: general AI and narrow AI (Broussard, 2018). Computer scientist and data journalist Meredith Broussard describes general AI as the "Hollywood kind of AI," like sentient robots or machines with consciousness. Narrow AI is what currently exists, which Broussard describes as a "mathematical method for prediction." Narrow AI is based on a model built by analyzing patterns in existing data sets (Broussard, 2018).

Narrow AI is what journalists are using right now. "Machine learning" is narrow AI, where the machine learns how to identify patterns and find what journalists (or the programmers) are looking for after analyzing hundreds of thousands of data points. Machine learning utilizes algorithms, which are "recipes" for a computer — the steps the algorithm takes to create a certain outcome ([Knight Center Courses], 2019). Another term often mentioned in conversation with AI is "deep learning," which can be seen as a subset to machine learning (M, 2018) or a type of algorithm of machine learning — others include decision tree, random forest and naive Bayes (Broussard, 2018). Deep learning refers to the deep neural networks, the layers of data, that the machines learn from just by being exposed to the data (M, 2018). For example, researchers used downloaded web images and their tags to train a model for image retrieval (Xu, Huang, & Wang, 2018). Right now, machine learning currently works by utilizing algorithms to

build a command structure for the machine to learn from. When a machine reaches the point when it can learn on its own without a command structure, that will be true machine learning, or general AI.

Artificial intelligence is currently used in many industries besides media — think self-driving cars and voice technology — but also in other lesser known ways such as in the military, where experts have been discussing development and ethical risks of lethal autonomous weapons (Hambling, 2018). In journalism, every aspect of the news production can be affected by machine learning and grouped into three main categories: story discovery, story production and story distribution. In general, the story discovery process includes using data to uncover leads, angles and topics for reporting, akin to data journalism as a field. Story production uses AI to help actually physically produce this story, such as writing and data visualization. As for story distribution, AI can be used as a tool to help with better reaching readers and distributing content to the organization's audience. AI can also be used in sub-steps like fact checking (Diakopoulos, 2018a), comment moderation (Etim, 2017), and advertising (Spangher, 2018), among others.

Advocates say artificial intelligence can save time and money, as it will liberate reporters from writing routine stories and allow them to devote more energy and time on larger and more complex topics (LeCompte, 2015). Machine learning can also help reporters analyze vast quantities of data faster than any human would and be able to distinguish patterns that would otherwise not be found (Broussard, 2015). There are still some tasks better suited to humans though, and those include situations where complex communication is needed, like dodgy sources and confidential information; expert thinking, such as contextualizing events; and dynamic environments, where news is not routine ([Knight Center Courses], 2019).

Media organizations have been keen on investing in artificial intelligence and are proactively experimenting with the technology. *The New York Times* is hiring more people skilled in artificial intelligence, machine learning, data science and mobile engineering (Slefo, 2018). In November 2018, the digital business magazine *Quartz* announced the launch of its AI Studio, which will partner with journalists at other media organizations to produce machine-learning based stories (Keefe, 2018).

Like all emerging technologies, though, there are numerous ethical issues that media organizations must address as they implement machine learning into their newsrooms. Dörr (2016) wrote about the need to be mindful about the kind of data used for machine learning in newsrooms, such as its origin, objectivity and collection methods. When searching for new stories, journalists also need to question leads generated by an algorithm (Diakopoulos, 2018b).

The use of algorithms in journalism also means that engineers and programmers will have a significantly bigger influence on how news is created and disseminated (Miroshnichenko, 2018). Journalists and programmers think differently and have different priorities when writing data-driven stories, and Linden (2017) suggests that the choice and evaluation of data should be a task for journalists while validation, standardization and normalization is normally a task for programmers.

Consideration of diversity among AI programmers is also critical. There are some notable examples of algorithms gone wrong, such as Microsoft's Twitter bot Tay, which tweeted racist content (Vincent, 2016) and Amazon's recruitment algorithm that was biased against women ("Amazon scrapped," 2018). Media organizations need to codify who will monitor algorithms when they malfunction, as Dörr and Hollnbuchner (2017) note that the role of the individual

diminishes while the importance of the media organization and the media system as moral agents increases.

Multiple studies have shown that readers are unable to distinguish between a human- or algorithm-written text (Miroshnichenko, 2018; Clerwall, 2014). Clerwell found that people rated human-written stories as more coherent, well-written, clear, less boring and more pleasant to read. Algorithm-written stories were more descriptive, informative and more boring but more accurate, trustworthy and objective, but there was no significant difference in the perception of the two texts except for the descriptor "pleasant to read" (Clerwall, 2014).

Another significant ethical issue is how authorship of algorithmic news should be attributed. Montal and Reich (2017) posed the question of whether bylines should be authorless, shared between machines and humans or attributed to just the algorithm, to the person buying the software or to the data provider. They found that people attributed authorship to the programmer or the organization, and most organizations didn't have a byline policy that differentiated between human and machine. They suggested that for algorithmic content generation, the byline should go to the software vendor or in-house programmer and should clearly detail the algorithmic nature of the piece and the methodology of using the data. For integrative content generation, such as when a human collaborates with an algorithm to produce content, they suggested the byline be attributed to the human journalist and clearly detail the pieces of the story generated by an algorithm as well as the algorithmic nature of the piece and the methodology.

Numerous studies have examined ethical dilemmas in integrating algorithms in aspects of the news production process, such as attributing authorship or transparency of the algorithm's usage. There are few studies that synthesize concerns media professionals should have for the

entire news production process, and no widely accepted conclusions on what should be considered for the media ecosystem as a whole. Therefore, this paper will explore overall ethical considerations media professionals should consider when implementing artificial intelligence in their newsrooms, and propose an ethical framework to guide their decision-making process.

How AI is Used in Journalism

Experts agree that usage of AI in journalism can be grouped into three steps of the news production process: story discovery, story production and story distribution. Using AI in story discovery can include using machine learning to scan documents and data for patterns that humans wouldn't normally be able to distinguish and creating bots to automatically alert a reporter or editor for something they specifically programmed the machine to take note of.

An example of this is computer scientist and data journalist Meredith Broussard's Story Discovery Engine (2015). Broussard programmed an algorithm to sort through data from the School District of Philadelphia to see if enough books were available for students to adequately prepare for state standardized tests. The software she created was a knowledge-based expert system that produced interactive data visualization as its output. This software could analyze more data sets than any single individual, and Broussard included examples of many of the engine's findings that could lead to potential stories.

Broussard's Story Discovery Engine is a good example of how AI can assist with investigative reporting. Other notable examples include the International Consortium of Investigative Journalists, which built a search engine to sort through massive amounts of documents like emails, PDFs, images and invoices (Romera & Gallego, 2018), and *BuzzFeed*, which published a story about training a machine to track flight patterns to identify hidden spy planes (Aldhous, 2017).

Media organizations have also programmed bots to detect news. One example is *The Intercept*'s ICE Watch Raids Twitter bot, which tweeted out descriptions of ICE raids in New York. Then there are algorithms like Duke University's Tech & Check Cooperative's ClaimBuster, which monitors claims in CNN transcripts that can be fact-checked, sending alerts

to reporters (Diakopoulos, 2018b). The challenging part about using algorithms to discover stories, though, is distinguishing what makes something newsworthy, as this is specific to each publication and its audience, among a variety of other factors (Diakopoulos, 2018b). This news judgment is something that, for the foreseeable future, will be hard to automate. The role of an editor who can make algorithm-informed judgments will be valuable, but editors also need to understand that what an algorithm spits out must be checked and validated — just like any suggestion or lead a human reporter pitches.

In terms of story production, the most prevalent form of AI is through automated content generation. Automated content writing (also known as "robot journalism" or "algorithmic news writing") can fall under the field of "computational journalism" and was first implemented in fields where clean, routine data is available, such as sports, finance, crime and weather (Miroshnichenko, 2018). The *Associated Press* began automating financial earning reports in 2014, using Wordsmith from Automated Insights. While outside engineers built the Wordsmith platforms, *AP* journalists write the story templates, says Lisa Gibbs, then business editor and now director of news partnerships at the *AP* (personal communication, March 12, 2019). Before using Automated Insights, the *AP* would publish 300 earnings report stories per quarter in 2014. Now the *AP* publishes 3,700 earning report stories a quarter (Peiser, 2019). This number will likely increase, both in terms of number of earning reports stories produced and number of companies reported on. *Bloomberg*, which also began heavily investing in automation in 2014, has in-house engineers who built their own templates to produce automated earnings reports stories (C. Eltzroth, personal communication, January 10, 2019). Both the *AP* and *Bloomberg* said that some stories automatically hit the wire without anyone approving, but the *AP* has the ability to turn off automation, Gibbs says, which they chose to do for banks at one point.

Bloomberg's stories — which are released to their Terminal customers, who rely on speed to make financial decisions — will have editors check a story and fill in more details if the machine didn't pick up enough, says Clay Eltzroth, product manager at *Bloomberg* (personal communication, January 10, 2019).

Finally, AI is also used in story distribution. Media organizations are creating bots and widgets to track reading history and recommend stories to users. For example, *The Washington Post*'s Clavis uses a term-frequency, inverse document frequency algorithm to "read" and categorize stories using keywords (Graff, 2015). It then does the same thing for its readers, and use the matching keywords to recommend stories to them (Graff, 2015). *Bloomberg* is also using AI in its story distribution process with the creation of The Bulletin for its mobile app, which generates one-sentence summaries of the most important stories in the *Bloomberg* network (Bloomberg, 2018). The summaries are refreshed each time the app is opened or manually refreshed.

Additionally, there are numerous academic researchers conducting studies on artificial intelligence in media. At the MIT Media Lab, open floors are filled with researchers tinkering with technology who are focusing less on "making algorithms faster" but rather the intersection of fields, like technology, media and society, says Mike Jiang, a research assistant in the viral communications group (personal communication, November 20, 2018). While Jiang says he usually doesn't work with AI in his research group, there are other projects in the lab, like Joy Buolamwini's study on biased AI systems in the Civic Media group (Buolamwini, n.d.). Buolamwini created the first algorithmic audit of gender and skin color in commercial facial analysis models.

How Media Organizations Integrate AI

U.S. news organizations large and small have started to integrate AI into their newsrooms, and the size of the organization, along with the type of content it produces, dictates how AI will be implemented. The *Associated Press* and *Bloomberg* have both invested heavily in automated content writing because their businesses depend on speed. Both organizations reach large, diverse audiences, with *Associated Press* content distributed to other print, video and audio media businesses and *Bloomberg* globally to business professionals and organizations. For earnings reports stories, for example, investors need to know information as quickly as possible. Gibbs noted the increase in earning reports stories the *AP* now publishes has actually impacted financial markets — companies that the *AP* has never written about before have been affected by more trading (Gibbs, 2019).

Both the *Associated Press* and *Bloomberg* are continuing to implement AI to help with routine tasks. The *AP* is developing a summary generating algorithm, Gibbs says (personal communication, March 12, 2019). When reporters file a story, they have to submit a summary so other media platforms, like radio stations, have the right story format for their audience. By training a machine to write the summaries, reporters can file faster and summaries will be more consistent. "Why should we be using our journalists, our skilled journalists, to essentially rewrite the same story five times? Is that really the best use of their time? I don't think so," Gibbs says (personal communication, March 12, 2019).

The *AP* works with startups to develop the AI software used in its newsrooms, but *Bloomberg* has its own in-house engineers. The purpose of using AI in its organization is not just for one task, but rather to implement across the company, Eltzroth says (personal communication, January 10, 2019). "It's the whole process of when something big happens," he

says. “What can you automate? What can you help the reporter do once they need to write the story? How can you help them in pulling content that’s associated with this event that occured in the past to help give insight into it, so it’s more than just the one thing that has happened?”

Bloomberg and the *AP* are both focused on producing breaking news stories, and while the same is true of *The New York Times*, its focus in using AI has centered more around other, less critical tasks to the news production process. *The New York Times* partnered with Google Jigsaw in 2016 to help with comment moderation. The algorithm flags and rates comments on the likelihood that would be approved or rejected. Bassey Etim, former community manager and editor at *The New York Times* who helped launch the project, says the algorithm helps moderators focus on the “middle section” of comments who have 40-60 percent likelihood of rejection (personal communication, October 8, 2018). This means the *Times* is able to keep more comments sections open at once. Previously, only a small handful of news and opinion stories would be open for comments at a time, but now every home-page story with a summary will have comments open, he says. The purpose of the algorithm is to help make the moderators’ jobs more efficient, but that isn’t the same as speed, Etim says. “Once you have a point of view about things that are fair to say and things that aren’t fair to say, then you are going to be able to create a product that reflects the organization, and create a product that makes people who are coming to your publication feel like they are still reading your publication,” Etim says. “We want *The New York Times* comments to feel like *The New York Times*.”

The New York Times has also started using Google Cloud to digitize its photo archive, which has even led to discovery of new information found on the backs of photos, like the time taken and location of the photo (Greenfield, 2018). Additionally, the *Times* also conducted an advertising project that crowdsourced readers’ emotions toward articles to help advertisers better

pair advertisements and stories (Spangher, 2018). While these are significant projects that help *The New York Times* as a whole, these AI projects aren't contributing to story discovery, story production or story distribution. This isn't a bad thing; one of the best qualities about AI is its number of applications in a variety of settings.

It's generally smaller newsrooms, like *Quartz* and investigative publication *The Intercept*, that are focused more on using AI to help with the reporting process and story discovery. These newsrooms, however, are still national. Few local U.S. newsrooms have yet to experiment with AI, and this could be due to a number of factors, particularly lack of time and resources. One notable local (yet still large) newspaper who has, though, is *The Atlanta-Journal Constitution*, whose journalists trained algorithms to sort through 100,000 documents looking for specific keywords. They then read this much smaller pile of 6,000 documents for an investigation on the state's doctors with sexual misconduct history who were still practicing (The Atlanta-Journal Constitution, 2016).

The Intercept — which has less than 50 employees — mostly uses AI as part of its document- and data-based reporting, says Rubina Fillion, director of audience engagement. Even then, AI is seen as more of a side project. *The Intercept* has mostly followed what other large organizations have done, as it doesn't have the in-house technology team or resources like other organizations have (R. Fillion, personal communication, March 22, 2018). *Quartz*'s AI Studio, which will pick six machine-learning based stories to work on with other news organizations this year, is focused on using AI as a reporting tool, says John Keefe, technical architect of bots and machine learning at *Quartz*. The studio is funded by a grant from the Knight Foundation, and *Quartz* plans to publish the methodology and process for these stories to use as examples for other journalists to use in their own reporting. *Quartz* is focused on working with small- to

medium-sized media organizations, as they are the ones who probably don't have the technology teams devoted to using AI (J. Keefe, personal communication, January 11, 2019). "We are going to document the process," Keefe says. "We are going write about what we do, not just for transparency sake, but more of the perspective of 'Here's what we did and here's how you can do this too.' So kind of teaching by example."

There is no standard way to use AI in the newsroom, which has allowed for creative and flexible projects to emerge. But while every news organization has pursued its own projects, a lack of widely accepted standards means that each organization is also following its own ethical guidelines. This could be problematic.

The Current Conversation Around Ethics

Media professionals interviewed have said that each of their newsrooms has their own discussions around ethical use of AI. There have been ethical discussion around each step of the story production process. Many of the potential ethical problems that face artificial intelligence also apply to the larger field of data journalism. One example is the importance of data quality — bad data means a bad story, Gibbs says (personal communication, March 12, 2019). "Whether it's an automated story or a human-produced story, to me a lot of these things that we are talking about are really kind of the same," she says. "It shouldn't really be any different. We are just using a different way of producing a story." When a data story is wrong, that's because the initial input data wasn't accurate. Just like using data in any story that's produced by a human journalist, data needs to be as unbiased as possible, cleaned and sorted to provide the most accurate story possible. This applies to AI in terms of the data used to train models as well as the data collected when finding stories.

Additionally, algorithms can't just be left alone once they are created — they need to be checked regularly, and updated if needed, depending on how the data set changes. For example, *The New York Times'* comment moderation algorithm is not very good at catching late-breaking political insults, Etim says. *The New York Times* doesn't allow political name calling in its comments, and if a commenter comes up with a new nickname for a politician, Etim says, the algorithm currently isn't going to catch it (B. Etim, personal communication, October 8, 2018).

Besides the issues that come with programming these algorithms, however, there are also ethical issues that come with how readers perceive work created by machines. Researchers have also found that readers are unable to distinguish between human-written and machine-written stories (Miroshnichenko, 2018; Clerwall, 2014). Miroshnichenko argues that in a regular Turing

test, there is an element of interaction as the human asks questions to the machine. In the journalistic Turing test, though, there is no interactive aspect but only the perception of a completed story. Style is the most difficult aspect of human speech for robots to overcome, he argues, but it doesn't matter — robots don't have to write better than humans but good enough.

However, it's not necessarily a bad thing that readers can't distinguish between human and robot written content — most of the time, reporters don't even want to write the content that is being automated anyways, Fillion says (personal communication, March 22, 2018). "It's not really valuable use of a human's time to produce those stories and it's taking them away from the kinds of projects they could be working on that really do need somebody who can dig deeper into something," she says. "So I think it's actually good to outsource some of those stories that are pretty formulaic to automation, and then be able to focus on things that can be a lot more in depth." Additionally, Tracy Ann Kosa, an adjunct professor at Seattle University School of Law, says that it's not bad at all that readers can't distinguish between human and machine written journalism. Instead, it's a testament to how advanced this technology is to manipulate and present a reality that someone else has created for us (T. Kosa, personal communication, March 29, 2018).

Instead, the "bad" aspect of this is the fact that there are no rules regarding usage of this technology and no conversations about the impact and implementation of it, Kosa says. All organizations that are using automated content generation have a line of disclosure at the end of all their stories, and all of them said this was an important part of the conversation from the start. For example, the *AP*'s reads, "Elements of this story were generated by Automated Insights (http://automatedinsights.com/ap) using data from Zacks Investment Research. Access a Zacks stock report on LULU at https://www.zacks.com/ap/LULU." Kosa likens the use of AI in

journalism similar to the public's initial perceptions of privacy, where it was hard to get people to understand the gravity and importance of it (personal communication, March 29, 2018). She says that privacy experts went about the situation poorly, because when incidents exploded they didn't give solutions to solve it. Instead, journalists need to explain why it's important to care about transparency, why the organization is being transparent and offer solutions that are manageable, implementable and reasonable.

Whether people care about transparency is irrelevant, Kosa says — it's a requirement for presenting any information in this context. That being said, Fillion mentions that including information on how AI or any technology or even just investigative reporting is used in a story, whether it be a sentence or a whole section on methodology, can also make readers more inclined to read the story. "I generally think that giving readers some behind the scenes information is valuable because it engenders more trust between the reader and the journalist," she says. "And also letting them into an entire process helps them actually understand how journalism is done, because there's a lot of misconceptions about that."

Being transparent is also important because of general trust issues in media, says Elite Truong, deputy editor of strategic initiatives at *The Washington Pos*t (personal communication, March 11, 2018). "If a reporter gets something wrong, it's easy to have that reporter be accountable," Truong says. "But when a machine gets something wrong, it's something you blame the institution on. It becomes faceless."

Legislation of AI is also a major debate. Gibbs notes that there is a natural objection reaction to the U.S. setting policies for media, but also pointed out that all these media organizations are running their own experiments and that eventually there needs to be some common understanding of best practices (personal communication, March 12, 2019). The *AP* has

a Standards Editor who answers any style or standards questions in the newsroom, and Gibbs says she worked with him to develop a standards entry for policy around disclosing automated content (personal communication, March 12, 2019). While the *AP* looked to professional bodies for guidance on usage of AI, in the future there may very well be a Stylebook section on automation, she says. This would be incredibly influential, as the AP Stylebook is used in nearly every newsroom in the United States.

Fillion also says that she is reluctant to trust legislation with the U.S.'s current political structure, pointing out that in Mark Zuckerberg's testimony for Facebook, legislators clearly didn't understand what the issue was or how algorithms based on the questions they asked. Potential legislation could be helpful but only if people know what they are doing, she says, so perhaps a subcommittee well-versed in AI and advised by experts (R. Fillion, personal communication, March 22, 2018). Kosa argues that in the debate over transparency and government censorship, organizations should still aim to be transparent first. She argues that what needs to be regulated are technologists as a whole, particularly in engineering. "I've really come increasingly to the conclusion that the problem isn't the technology itself," she says. "The problem is the absence of any accountability or social responsibility in the profession, and in the development of the technology. It's far too impactful on the rest of the world." Other professions, like doctors, lawyers and architects, all understand that their decisions have an impact on society, but technologists haven't reached this realization yet and are still building and having a disruptive and massive impact on society.

Additionally, using AI in journalism means a massive amount of data collection. Something news organizations need to be aware of is that while there are plenty of free tools that newsrooms can use, especially if they don't have resources to invest in their own data team, they

need to be clear on what data is being collected, Truong says. “Those are your loyal readers, your subscribers, your audience,” she says. “If you choose to use those free tools you just want to be very careful to understand what it means to opt in your users and your viewers.”

Building an Ethical Framework

For general media ethics, Dörr notes that there are three spheres: organizational, professional and social. The organizational sphere represents the ethics of the media system, the professional sphere encompasses individual ethics and the social sphere includes ethics of the audience (Dörr, 2016). Bryson (1998) argues that an ethics system needs to be self-regulating, and that humans tend to put language, mathematics and "reason" as the key characteristics of human life, meaning that we tend to over-identify with machines, causing us to elevate the worth of the machine and lower the worth of humans. Bryson argues that AI can be used to rationalize human ethical decisions.

The most important idea for ethical use of all technology, regardless of field, is having a human involved, otherwise known as "human in the loop" systems, Kosa says (personal communication, March 29, 2018). "I think it's less about journalists themselves taking over or being taken over by some sort of technology, whether that's automated content writing or something else," Kosa says. "And it's more about the recognition that even the very best of the best of technology X, automated content writing for the purposes of our conversation, still necessitate having a person be involved before it's unleashed anywhere."

Journalism has traditionally been guided by self-developed, monitored codes of ethics in each individual media organization. In the industry as a whole, though, there are certain ethical standards that journalists are expected to meet. There have been frameworks built for journalists already, like the Society of Professional Journalists, whose Code of Ethics names four principles and includes practical and constructive rules for journalists to follow when reporting. Mark Pearson applied Buddhism's Noble Eightfold Path to journalism to a create a practice he deemed "mindful journalism." Experts have also built frameworks for general use of AI, such as the

European High-Level Expert Group on Artificial Intelligence, and OpenAI, an organization focused on discovering safe use of artificial general intelligence. But based on conversations conducted for this report, it appears that none of the organizations that have integrated AI in their newsrooms have added a section to their code of ethics that addresses artificial intelligence.

To formulate an ethical framework for media organization usage of artificial intelligence, different guiding prompts were organized by each step of the news process: story discovery, story production and story distribution. This framework is purposefully not an ethical code, as ethical codes are something each individual news organization needs to build for itself. As the technology is evolving so quickly, this framework is meant to be prescriptive.

General Guidance

- Regularly maintain standards for data collection, usage and distribution currently in place in the media organization.
    - Media organizations should have regular meetings and/or discussions to reflect on how data is being used both for editorial and business purposes. These meetings are especially critical when a new technology or process is being implemented, so frequency of meetings will likely increase.
- Articulate these standards in the organization's code of ethics.
    - If they are not already articulated, they need to be. Use of data in journalism has been around long enough and will be clearly become a vital and critical aspect to producing news that it must become standardized for everyone in the organization, not just specific data journalists or teams.
- Document the process for creation of every algorithm used by the media organization.

    - This is especially critical for legal reasons, as this documentation can be used as defense for any sort of serious claims, such as someone claiming an algorithm created libel.
- Publish public-facing methodologies for producing algorithms that are easily accessible and understandable for the public.
    - Transparency is absolutely key and there is no excuse not to provide it. It doesn't matter if people don't read it — it needs to be readily available for public consumption. It also can't just be the lines of code or computer language jargon. If an average reader of the media publication wanted to know how the algorithm was created, they should be able to understand the general process without having a technical background. As many details as possible should be released to the public as long as they don't jeopardize privacy and/or competitive advantages.
- Frequently analyze models for bias. Who is analyzing them?
    - There should be a standard set of questions that reporters and editors can use to identify bias in a model. Additionally, when developing the algorithms and discussing the potential biases that might become encoded, there should be as diverse a group of people included as possible. This includes race, gender, income level, job position and any other characteristics that could be pertinent to a specific story.
- Make clear to both engineers and journalists what editorial and ethical standards are for algorithm building.
    - The relationship between engineers and journalists is obviously different in each organization, but regardless each party should always be clearly communicating

what expectations are in terms of how these algorithms are coded, especially for automated content generation. Engineers and computer scientists have greater control over how news is written than ever before, and they don't necessarily have the same news training as traditionally-trained journalists.

Story Discovery

- Periodically evaluate data being used to train algorithmic models.
    - Training data can become outdated quickly. For every algorithm that a media organization uses to find story leads, a person should be assigned for regularly checking on the data the model uses to generate suggestions. The frequency of these evaluations depends on the length of the project, so there would be fewer evaluations if this algorithm was used for a six-month project compared to a three-year project. It also matters more if these algorithms are used in perpetuity for story lead generation.
- Create a process to evaluate claims, leads and outputs that an algorithm generates.
    - There needs to be someone to fact-check and critically examine any serious, pursuable claims an algorithm generates before reporters begin work.

Story Production

- Regularly examine an algorithm for context.
    - If an algorithm is being used to automatically generate content, reporters and editors also need to keep in mind the context of the environment the algorithm is reporting on. For example, if a law changed or a financial market was hit abruptly, that could lead the algorithm to write stories that are out of context. In

these types of situations, teams should discuss turning the algorithm off while it is improved.

Story Distribution

- Understand what data third parties are collecting if using a third-party algorithm or tool.
    - Editors and reporters need to weigh the cost of using these algorithms and critically evaluate if it is worth giving up private and potentially competitive data about its audience to a third-party. Media organizations should feel protective of this data.
- Readers should understand how and in which parts algorithms were used in any part of reporting or idea generating process.
    - Readers should not feel confused as to how a machine was used in the creation of a story. Particularly if an algorithm was used to write a story, attribution should be given to the algorithm and in the disclosure at the end of the story there should be a link to more information on how the story was generated and who was responsible for creating the algorithm.

Conclusion

As with all emerging technologies, the excitement and potential for artificial intelligence to change journalism needs to be balanced with thoughtful, ethical decision making regarding how this technology is implemented. Right now, most newsrooms using AI are all national publications, even though they vary in size. AI might be limited to national publications for right now, but soon enough it will be accessible to smaller, local newsrooms as well. “The trick of making something successful is being able to give that service to people cheaply,” Eltzroth says. There are a lot of potentials for using local data as well, and journalism organizations are getting increasing funded to explore these opportunities. The *AP* recently received a grant to distribute data journalism to local news organizations, Gibbs says.

Someone that is on everyone’s minds, though, is the future of jobs in journalism. There is no doubt that the industry is still facing job loss, but there is also a need for people who can understand and manage emerging technology like AI. More organizations are now offering jobs that focus on being liaisons between journalists and engineers, which is similar to the jobs Fillion, Truong, Gibbs, Eltzroth and Keefe all have right now. In fact, Fillion notes that it’s hard to recruit people who have skills both as a journalist and a technologist because they often take jobs in tech, which pays better.

As for future journalists, everyone interviewed agreed that not everyone needs to know how to code. Keefe says that people who know how to code will probably be able to find more stories than journalists who don’t, but everyone will always have their own specialty — which is how it’s always been before. What’s important is data literacy, something both Truong and Gibbs stressed. “There is still a huge vacuum in basic data literacy, I mean I think there are way too many journalists coming out of journalism school not understanding how to use data on their

beat," Gibbs says. It's knowing how to interpret data, and also being able to talk to people who do code, that is important and lacking, Fillion says.

As for limitations to this research, I only talked to 12 people. Some were from the same news organizations and not all of them were specifically mentioned in this research. Some of these people guided my thinking and research, but didn't answer specific questions about their newsrooms. There are numerous other organizations that are also using artificial intelligence that I wish I could have spoken to, such as *ProPublica, The Wall Street Journal* and *The Los Angeles Times*. Additionally, this report only looked at U.S. newsrooms and media organizations, but there are numerous European and Asian media organizations experimenting with AI. I didn't ask the same set of questions to each person and my questions improved over time. This thesis was a year-long process and I had to overcome a relatively steep learning curve — one year ago, I knew almost nothing about artificial intelligence. I had no idea how it was even used or created in the first place, so just understanding what the technology's strengths and limitations were took time. This meant that what I thought were ethical issues when I began researching changed and grew far more nuanced by the end of writing this report.

I see the most potential in using artificial intelligence as a tool for generating story leads and reaching readers. Automated content generation is the focus right now, but once the shock of having machines writing stories passes (and it already is in this process), people will realize the flexibility of the technology means it can aid humans to help them produce better stories, be more efficiently and most importantly, realize the technology will not take their jobs. Once AI becomes accessible enough to local media organizations this will be especially clear. Using AI as a tool to help on the business side of a media organization — such as how *The New York Times*

is using it to gauge emotional resonance of a story and place advertisements — could have the potential to reshape parts of the business model.

References

Aldhous, P. (2017, August 7). We Trained A Computer To Search For Hidden Spy Planes. This Is What It Found. *BuzzFeed News*. Retrieved April 11, 2019, from https://www.buzzfeednews.com/article/peteraldhous/hidden-spy-planes

Amazon scrapped 'sexist AI' tool. (2018, October 10). *BBC*. Retrieved April 11, 2019, from https://www.bbc.com/news/technology-45809919

Bloomberg. (2018, September 18). *Bloomberg Media's Innovation Lab launches "The Bulletin" – An AI-Powered News Feed for Bloomberg Mobile App Users* [Press release]. Retrieved April 11, 2019, from https://www.bloombergmedia.com/press/bloomberg-medias-innovation-lab-launches-bulletin/

Broussard, M. (2015). Artificial Intelligence for Investigative Reporting. *Digital Journalism, 3*(6), 814-831. April 11, 2019, http://dx.doi.org/10.1080/21670811.2014.985497

Broussard, M. (2018). *Artificial Unintelligence: How Computers Misunderstand the World* (1st ed.). Cambridge, MA: MIT Press.

Bryson, J., & Kime, P.. (1998). Just another artifact: Ethics and the empirical experience of AI presented at Fifteenth International Congress on Cybernetics, Edinburgh, January 1998.

Buolamwini, J. (n.d.). Actionable Auditing: Coordinated bias disclosure study. Retrieved April 11, 2019, from https://www.media.mit.edu/projects/actionable-auditing-coordinated-bias-disclosure-study/overview/

Clerwall, C. (2014). Enter the Robot Journalist. *Journalism Practice, 8*(5), 519-531. Retrieved April 11, 2019, from https://www.tandfonline.com/doi/abs/10.1080/17512786.2014.883116

Diakopoulos, N. (2018, June 28a). There are a lot of rote tasks a good AI interviewer could do for you. *Columbia Journalism Review*. Retrieved April 11, 2019, from https://www.cjr.org/tow_center/artificial-intelligence-reporting-interviews.php

Diakopoulos, N. (2018, November 28b). An algorithmic nose for news. *Columbia Journalism Review*. Retrieved April 11, 2019, from https://www.cjr.org/tow_center/an-algorithmic-nose-for-news.php

Dörr, K. N. (2016). Mapping the field of Algorithmic Journalism. *Digital Journalism, 4*(6), 700-722. April 11, 2019, http://dx.doi.org/10.1080/21670811.2015.1096748

Dörr, K. N., & Hollnbuchner, K. (2017). Ethical Challenges of Algorithmic Journalism. *Digital Journalism, 5*(4), 404-419. Retrieved April 11, 2019, from https://www.tandfonline.com/doi/full/10.1080/21670811.2016.1167612

Etim, B. (2017, June 13). The Times Sharply Increases Articles Open for Comments, Using Google's Technology. *The New York Times*. Retrieved April 11, 2019, from https://www.nytimes.com/2017/06/13/insider/have-a-comment-leave-a-comment.html

Gibbs, L. (2019, March 13). How AI Has and Will Impact the News Media Industry. In *SXSW 2019, Austin, Texas*.

Graff, R. (2015, June 3). How the Washington Post used data and natural language processing to get people to read more news. *Knight Lab*. Retrieved April 11, 2019, from https://knightlab.northwestern.edu/2015/06/03/how-the-washington-posts-clavis-tool-helps-to-make-news-personal/

Greenfield, S. (2018, November 9). Picture what the cloud can do: How the New York Times is

using Google Cloud to find untold stories in millions of archived photos [Web log post]. Retrieved April 11, 2019, from https://cloud.google.com/blog/products/ai-machine-learning/how-the-new-york-times-is-using-google-cloud-to-find-untold-stories-in-millions-of-archived-photos

Hambling, D. (2018, September 14). Why the U.S. Is Backing Killer Robots. *Popular Mechanics*. Retrieved April 11, 2019, from https://www.popularmechanics.com/military/research/a23133118/us-ai-robots-warfare/

Keefe, J. (2018, November 20). Announcing the Quartz AI Studio, designed to help journalists use machine learning. *Quartz*. Retrieved April 11, 2019, from https://qz.com/1464390/announcing-the-quartz-ai-studio-designed-to-help-journalists-us-machine-learning/

[Knight Center Courses]. (2019, January 8). An Overview of Algorithmic News Media [Video File]. Retrieved from https://www.youtube.com/watch?v=gtGGh4pRc1Q

LeCompte, C. (2015, September 1). Automation in the Newsroom. *Nieman Reports*. Retrieved April 11, 2019, from https://niemanreports.org/articles/automation-in-the-newsroom/

Linden, C. (2017). Decades of Automation in the Newsroom. *Digital Journalism, 5*(2), 123-140. Retrieved April 11, 2019, from https://www.tandfonline.com/doi/full/10.1080/21670811.2016.1160791

M, V. (2018, May 7). Artificial Intelligence vs. Machine Learning vs. Deep Learning. Retrieved April 11, 2019, from https://www.datasciencecentral.com/profiles/blogs/artificial-intelligence-vs-machine-leaning-vs-deep-learning

Miroshnichenko, A. (2018). AI to Bypass Creativity. Will Robots Replace Journalists? (The Answer Is "Yes"). *Information, 9*(7). Retrieved April 11, 2019, from https://www.mdpi.com/2078-2489/9/7/183

Montal, T., & Reich, Z. (2017). I, Robot. You, Journalist. Who is the Author? *Digital Journalism, 5*(7), 829-849. Retrieved April 11, 2019, from https://doi.org/10.1080/21670811.2016.1209083

Peiser, J. (2019, February 5). The Rise of the Robot Reporter. *The New York Times*, p. B1. Retrieved April 11, 2019, from https://www.nytimes.com/2019/02/05/business/media/artificial-intelligence-journalism-robots.html

Romera, P., & Gallego, C. S. (2018, July 3). How ICIJ deals with massive data leaks like the Panama Papers and Paradise Papers [Web log post]. Retrieved April 11, 2019, from https://www.icij.org/blog/2018/07/how-icij-deals-with-massive-data-leaks-like-the-panama-papers-and-paradise-papers/

Society of Professional Journalists. (n.d.). Code of Ethics. Retrieved April 11, 2019 from https://www.spj.org/pdf/spj-code-of-ethics.pdf

Spangher, A. (2018, October 31). How Does This Article Make You Feel? *Times Open on Medium*. Retrieved April 11, 2019, from https://open.nytimes.com/how-does-this-article-make-you-feel-4684e5e9c47

Slefo, G. (2018, December 3). New York Times Plans to Invest Heavily in AI to Improve Personalization. *AdAge*. Retrieved April 11, 2019, from https://adage.com/article/digital/york-times-poised-copy-facebook/315831

The Atlanta-Journal Constitution. (2016). Doctors and Sex Abuse. Retrieved April 11, 2019, from http://doctors.ajc.com/

Vincent, J. (2016, March 24). Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day. *The Verge*. Retrieved April 11, 2019, from https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist

Webb, A., Giralt, E., Levkowitz, R., Palatucci, M., & Perez, K. (2019). *2019 Trend Report For Journalism, Media & Technology* (Rep.). Retrieved April 11, 2019, from Future Today Institute website: https://futuretodayinstitute.com/2019-journalism-media-tech-trends/

Xu, H., Huang, C., & Wang, D. (2018). Enhancing semantic image retrieval with limited labeled examples via deep learning. *Knowledge-Based System, 163*, 252-266. Retrieved April 11, 2019, from https://www.sciencedirect.com/science/article/pii/S0950705118304295?via=ihub