CONCORDIA UNIVERSITY

School of Graduate Studies

This is to certify that the essay prepared

By: Elia Tanios El Habre

Entitled: A Case For Switching From Self-Fulfilling Prophecies To Rights In AI Journalism

and submitted in partial fulfillment of the requirements for the degree of

Master of Arts (Digital Innovation in Journalism)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final Examining Committee:

_____ Chair
David Secko

_____ Examiner
David Secko

_____ Examiner
Aphrodite Salas

_____ Supervisor
David Secko

Approved by _____

Chair of Department or Graduate Program Director

_____ _____

Dean of Faculty

CONCORDIA UNIVERSITY

A Case For Switching From Self-Fulfilling Prophecies To Rights In AI Journalism

by

ELIA TANIOS EL HABRE

An Essay

submitted in Partial Fulfillment of the Requirements

for the Degree of Masters of Arts (Digital Innovation in Journalism)

to the Department of Journalism

of the Faculty of Arts and Sciences

at Concordia University

Montreal, Quebec, Canada

May 2020

ABSTRACT

Journalists have mostly focused on AI news from a business lens, offering stakeholders reports on software utilities and their potential disruption. This application-based approach has left a gap with end-users, including journalists themselves, in questioning whether AI solutions are always desirable, and when so, whether citizen rights to data ownership and privacy are respected in the process. Following major security breaches and calls for big data regulation, this topic is compelling for journalistic examination as some citizens show reluctance over sharing personal information, even during the 2020 COVID-19 pandemic. In that context, very little academic research exists on how journalists cover AI news and therefore few guidelines exist on its best practices. In this essay, I summarize existing reports on how AI is covered, delineate what makes this news beat unique, and how journalists can co-opt rights-based approaches to engage their audience. Seven alternative angles are discussed in combination with practice advice for moving the journalism conversation on AI forward. This approach seeks to encourage further journalism studies research into the topic.

CONTENTS

1. INTRODUCTION

Journalists have traditionally covered Artificial Intelligence (AI) stories under the bigger umbrella of technology and business, thereby raising considerable limitations. Brennen's analysis of AI news in the UK found that journalists tended to amplify AI's implications by focusing on their potential rather than their current capabilities. This tendency is not surprising given that 60% of news analyzed by Brennen centered on an industry product, and that 33% of sources cited were industry-related (Brennen, 2018). An even higher percentage was scored in US newspapers which cited industry-associated individuals or businesses up to 65% of the time. (Chuan, Tsai, & Cho, 2019). In agreement with Brennen, Haider & Gandomi's analysis of how big data is covered concludes that not only do most news topics revolve around startups and the digital economy, but also that many news outlets heavily rely on press releases for their daily technology stories (Haider & Gandomi, 2019). Haider & Gandomi note that "[n]ews stories were consistent in singing the praise of the same technologies, geographies, and processes. The coverage seems repetitive and uncritical." Chyi & Lee refer to this cheerleading effect as "churnalism", where news and promotional material are often combined, analyzing examples of how Apple products are overrepresented in articles compared to their competition, and often in a positive light (Chyi & Lee, 2018). This Essay is couched in these issues and seeks to provide a perspective on ways forward.

More generally, many journalists tend to hardly question whether technological solutions are the best given the problems at hand (Brennen, 2018). Data journalist Meredith

Broussard offers a defying stance towards what she calls “technochauvinism”, or the supposition that tech solutions are always better (Chen, 2018). As most readers can extrapolate from their daily lives, software updates can always backfire leaving loopholes behind as inevitable lures for future exploitation. Tinkering with tech works both ways, to improve or trick an algorithm, and depends on the intent and skill of the humans behind it. In fact, AI Textfooler, a tool that changes words into synonyms, was able to reduce the accuracy of a movie reviews’ classification AI (to sort positive from negative) from 90 to 10% (Knight, 2020b). In short, AIs can fool other AIs. Tech advances are not necessarily a solution, but a tool that has the potential to augment existing flaws. However, in such a technologically positive framework, journalists can also act as counterbalancing “analysts” to contextualize the underlying social causes and effects behind AI advances (Hanitzsch & Vos, 2018).

While this analytical role is worthy of introspection, it is complicated by stakeholders who hold varying priorities, including AI communities themselves. Despite evidence opposing “the claim of a bias toward negative sentiment in the news media coverage of AI”, several communities are still antagonistic toward journalists and have argued against media presence at AI conferences (Garvey & Maskal, 2019). Many journalists also conflate AI with automation and liken changing work environments to previous tech revolutions that came with computers and the internet (Broussard et al., 2019). Journalists and AI communities are thereby entangled for better or worse, which for some topics may not be overly important, but for a technology that aims to disrupt everything from Netflix recommendations to the future of work, while mostly being developed behind the confidentiality of proprietary information, this is an issue

that requires journalists to be fully immersed. This can be achieved not only by understanding the technology, to enhance their questioning, but also covering all stakeholders in the field, not just its developers.

There's plenty of academic discussions on the effect of AI on journalism, such as how it affects journalists and editors (Hansen, Roca-Sales, Keegan, & King, 2017; Ali & Hassoun, 2019; Biswal & Gouda, 2020). This ranges from applications in document classification, language analysis, lead generation, and the monitoring of breaking news. In some cases, journalists have been invited to explore ways to be design partners on AI solutions that would enable them to better research and verify information (Gutierrez-Lopez et al., 2019). Jonathan Stray describes how AI could help investigative journalists in story production through "data wrangling". By potentially leveraging the power of "data extraction from documents and cross-database record linkage", AI could shave off time usually reserved for tedious tasks within a data-driven investigative story (Stray, 2019). Beyond story production, AI also contributes to news curation. The news experience isn't about serendipity anymore, as our media ecosystem can be dominated by Google news and other news curation services, which sort and prioritize headlines based on the temporality and the authority of stories among others. As described by Nechushtai & Lewis, "while there is broad recognition of what machines can do, there is growing unease about what they should be doing", raising questions of how this curation encroaches on the gatekeeping role of journalists (Nechushtai & Lewis, 2019; Hanitzsch & Vos, 2018).

However, even as academic discussions touch on what journalists want from AI, very little is found on how journalists should be covering AI stories. This in turn is not enough to help inform, guide and empower citizens on questioning what society at large wants from AI. This discussion is time sensitive, especially as the use of facial recognition AI becomes more popular with government and private entities. In early 2020, a moratorium on the use of facial recognition by the US government was proposed in a letter drafted by the Electronic Privacy Information Center (EPIC). 40 signatory groups, ranging from advocacy to press organizations, have endorsed this letter recommending a ban pending "further review". In short, organizations were fearful that the government's current use of facial recognition was dissonant with consumers' concerns. The letter's authors end by writing: "[t]he rapid and unregulated deployment of facial recognition poses a direct threat to "the precious liberties that are vital to our way of life"" (Chen, 2020a).

This lack of clarity on what governments and societies want from tech companies is not a new discussion in most democracies. In many cases, it has resulted in consumers and even tech employees asking their companies to stop dealing with the government, especially on military contracts (Campbell, 2018). This has fueled many in the tech world to call for the need for elected governments to police tech, and not the other way around. As Carol Ann Browne from Microsoft puts it, our faces deserve the same protection as our phones, and facial recognition should be regulated. She adds: "Democracy has always depended on the ability of people to meet and talk with each other, and share their views both in private and in public. This relies on people being able to move freely and without constant government surveillance"

(Smith & Browne, 2019). It is interesting that thought leaders are taking it upon themselves to pen articles, calling for their own regulation. This could be a golden opportunity for international tech journalists to call for seizing the moment. After all, tech consumers are not just North Americans; so, if compromises are willing to be made to better protect user privacy, why not coalesce this will and enhance the experience for users worldwide.

This essay argues that journalists should cover AI primarily as a data-powered technology, and that data privacy is a right to which citizens are entitled. This has been highlighted by the Privacy Commissioner of Canada Office, the Electronic Frontier Foundation, and other stakeholders and journalists (Therrien, 2020; O'Brien, 2020; Raval, 2019; Menand, 2018). Citizens should have a say in big data governance. By acting as mediators between the public and opaque algorithms, journalists can preserve their roles as sense makers and authenticators of news. These roles, as previously described by Hanitzsch & Vos, fall under a "critical-monitorial function" of journalism and can range from passive monitoring to an active watchdog vis a vis institutions and businesses (Hanitzsch & Vos, 2018). Journalists can also potentially act as mobilizers, and encourage audiences to engage in the technological conversation. They would discourage themselves and fellow citizens from sleepwalking through decision making and direction setting about advancements in the industry, a termed referred to as "technological somnambulism" (Langdon, 1986). To achieve this, the essay will highlight how AI is dependent on big data, how advances in the field are currently covered, and how the conversation can be shifted from covering applications in their social context to questioning the cost of AI as an answer.

2. AI AND DATA

According to IEEE-USA, "Artificial Intelligence (AI) is the theory and development of computer systems that are able to perform tasks which normally require human intelligence such as, visual perception, speech recognition, learning, decision-making, and natural language processing" (IEEE Board of Directors, 2019). Most AI practitioners use machine learning to achieve this goal (Vincent, J., 2019; Liu, 2020) and agree that "[i]t's no longer a primary objective for most to get to AI that operates just like a human brain, but to use its unique capabilities to enhance our world" (Marr, 2018). In Artificial Unintelligence, data journalist Meredith Broussard explains that machine learning allows a system to "become more accurate at performing a single, specific task according to a specific metric that a person has defined" (Broussard, IEEE Xplore, & MIT Press, 2018). Artificial intelligence is, however, not just a field of study in computer science where programmers code for algorithms. Unlike traditional coding, machine learning trains algorithms to forgo specific sets of instructions and infer from patterns to power its predictions. For example, recognizing the picture of a car doesn't depend on a programmer outlining inclusion/exclusion criteria (needs to have a door, cannot have zero windows, etc.) but on a large data set of cars that are labeled as such. From this starting point, the algorithm may posit that specific height to width ratios, color combinations and other variables are reliable patterns, each with its own relative importance. As Broussard puts it "[t]he machine is provided with the training data and labeled outputs. We essentially tell the machine what we want to find, then fine-tune the model until we get the machine to predict what we know to be true" (Broussard et al., 2018) Each time the algorithm has an incorrect

prediction, the weight previously assigned to each variable changes to accommodate for a higher success rate. As such, the machine-learning powered AI relies more on troves of data, the more varied the better, and less on refined code.

AI's power is a question of goal definition, computing power and labeled data. While the goal and computing power are mostly set by programmers, this doesn't necessarily apply to data. The data training sets used could be publicly available from natural systems such as constellations, seismic patterns, etc., or built from people's activity. In the latter case, when data pertains to specific groups of people, the question is one of ownership. Social media users inevitably create labeled data sets, from assigning specific hashtags to their tweets, locations to their restaurant reviews, and names to their photo tags. These may seem trivial on an individual basis, but on a large scale it has given leverage power to social media companies that have made their fortunes by compiling digital identities of their users. As such, AI use of publicly or privately shared data about citizens is a matter of privacy and ownership rights and should be treated as such by journalists.

Current journalism practice, however, often ignores human rights and other conditions affected by AI advances beyond automation and the job market. Theoretically, one could view journalism on AI as gaining some influence from current legal discussions. On one hand, the co-director of the Berkeley Center for Law & Technology Sonia Katyal argues that the "impending conflict between the protection of civil rights and artificial intelligence" can only be remedied by holding private companies accountable for their algorithms (Katyal, 2019). Additionally, co-

director of the International Human Rights Program, Petra Molnar highlights that "[t]echnology is not inherently democratic" and makes the case to examine AI use in human migration management (Molnar, 2019). In like manner, almost half of the tech experts who responded to a recent Pew Research Center query express pessimism on the effect of technology on democracy (Anderson & Rainie, 2020). One respondent stated that "[t]ech and AI are owned by their creators, the top 1%, with decisions made about the 100% in every sector of society that have little to no transparency, human judgment or much recourse", which implies a big shift in how society deals with democracy is upon us. This also puts the onus on journalists, who've not only had their business model hurt by tech platforms, but also their voices drowned by the overwhelming volume of fake news, to step up. As cybersecurity journalists J.M Porup opines on his beat, he cautions that "[l]ow-tech journalists reporting on these issues to low-tech audiences often confuse the issue […] The outlook is grim, and without more tech-savvy journalists to raise the alarm, I am pessimistic about the future of our political liberty" (Anderson & Rainie, 2020). On the other hand, media studies researchers tend to focus less on how to cover AI and tech, and more on the use of big data itself to produce journalism, and its implications. For example, Fairfield & Shtein posit that "[b]ig data technologies increase costs of compliance with traditional ethical values, while steeply lowering costs of invasion of privacy". They make the case for an ethical framework that is both stable and flexible for social scientists and journalists in this field (Fairfield & Shtein, 2014). Similarly, Lewis & Westlund not only focus on ethics but additional conceptual lenses such as epistemology and economy to "explore both contemporary and potential applications of big data for the professional logic and industrial production of journalism" (Lewis & Westlund, 2015). By exploring "the intersecting norms of

journalism and big data" in journalistic endeavors such as publishing large data sets and obtaining consent, not much is said on the adequacy of tech approaches for answering questions of public concern on data privacy and tech. As such, many academic discussions center around the aftermath of the use of big data and AI, which translates into a tech-assumptive approach spilling into current journalistic practices, which are explored in the next section.

3. CURRENT JOURNALISM PRACTICES ARE APPLICATION-FOCUSED

As mentioned in the introduction, most journalism coverage of AI advances centers around its products and applications. A recent review of newspapers published in the Netherlands in the last decade shows that AI coverage ranged from robots with football capacities to chatbots and fake news (Vergeer, 2020). In all cases though, AI's capabilities are mainly dependent on pattern recognition in big data. This is often a myopic overview that might lead to problematic definitions of a problem or issue. In fact, big datasets are currently driving the largest educational experiment in the world. Squirrel, a Chinese AI tutoring program, has over a million students enrolled since it was founded five years ago. It has achieved unicorn status (1B$ valuation) by promising to remedy educational deficits in an already highly-competitive academic scene. Squirrel evaluates students based on academic objectives and dispenses exercises catered to what they need to learn and practice (Hao, 2019). This is not necessarily what they want to learn, or something that encourages pursuits or curiosity outside

specific problem areas. It is basically crafting high achieving test takers, which begs the question of whether this is the goal of education. AI cannot answer that question. But Squirrel will strive to achieve the goals it was prescribed. Hao doesn't delve much into the causes that may lead specific students to underachieve or interrogate parents or school boards on whether Squirrel's goals align with theirs. As Broussard describes her own experience with the school system, she reminds readers that "[e]ducational standards are not natural laws. They are ideas that arise from specific political and ideological contexts" and cautions on how the rush to enforce an engineering solution onto education might "obscure[s] the social issues that are at stake" (Broussard et al., 2018).

Journalists are implicated when social issues are obscured. Not only can big data drive ambitious AI solutions, but it could also mislead journalists covering that beat as to what the underlying story might actually be. For example, in September 2019, the world press was covering a court decision after a UK citizen sued to stop facial recognition use by the New South Wales police. The man lost in a city that is notoriously busy with CCTV cameras: London has 48.4 for every 1000 citizen, higher than Beijing. However, most of the New York Times articles on that topic focused on the country's lax attitude towards sacrificing privacy for better security (Satariano, 2019). However, adoption of AI in policing was driven by cost cutting, something that Satariano only mentions in one line towards the end. It was not the case that AI deployment lead to fewer officers on the street, but the other way around. The technology was not adopted because it was better than having officers on the ground, but to remedy that shortage in the first place. The reporter doesn't question neither what lead to the funding cuts,

nor whether the streets were safer with more police officers on the ground. Satariano instead focuses on the application at hand, the use of facial recognition, without providing much context on whether this application is justified. As such, the reader is asked to accept that surveillance is the most cost-effective solution. The frame of privacy, security and number of arrests is always appealing, but could also benefit from a contextual view of effective policing. Big data is driving AI decisions, but it doesn't have to dominate journalistic framing.

4. UNIQUE ISSUES TO AI

In addition to the need for social context when reporting on technological advancements (see section 3), some issues are unique to the AI field and should arguably be ever-present in its coverage. Recognition of six unique issues could lead to clearer perception on what sets AI apart from other tech, which can also be utilized to contextualize it in broader social issues outside of the tech beat.

4.1. Communicating uncertainty

First, setting goals for machine learning to train an AI system is exceptionally problematic. As Dr. Stuart Russell puts it, an intelligent AI system is set to achieve its goals, while we should hope to build beneficial AI that is set to achieve our goals. He adds that "[m]achines that have our objectives as their only guiding principle will be necessarily uncertain about what these objectives are, because they are in us — all eight billion of us, in all our

glorious variety, and in generations yet unborn — not in the machines" (Russell, 2019) Russell is a proponent of regulation and careful consideration on how goals are set in AI systems, in any area from education to self-driving cars, as we still don't know what we want from these systems, and what our priorities are in them.

Dr. Russell's opinion is unpopular not only because it's calling for regulation in a new booming field, but also because science journalists already walk a tight rope when communicating uncertainty. In chapter one of Communicating Uncertainty, Stephen Zehr highlights that "[i]ndeterminacy involved situations in which not all the parameters of a system and their interactions are fully known. Ignorance refers to situations in which it is not known what is not known" (Friedman, Dunwoody, & Rogers, 1999). Uncertainty in AI goals is not a matter of ignorance but one of indeterminacy. The reason for that uncertainty in AI goals needs to be highlighted when AI applications are reported on. In fact, Russell's call could be interpreted as a middle ground between a laissez faire approach and helicopter parenting. In the first case, it's similar to teaching a child how to eat based on hours of watching people eat. That very wide goal definition might lead to the conclusion that eating with one's hands gets the job done, and maybe the most food in, without considering that it's not a common custom in that child's family. On the other hand, a narrower goal definition might include adding multiple restrictions, and maybe coding in "holding fork with right hand" while that child is left-handed. In both cases, the objective is either way too loose or too tight, which might entail the need to set smaller goals (eating bites). Journalists ought to be capable of discerning ambitious AI goals, and explaining how prudent smaller goals do not signal ignorance.

4.2. Murky ethical spaces

A second issue that is unique to AI is determining responsibility when a system fails. This need exists because in many AI cases, content creation is not always correlated with liability. For example, an AI can't be an inventor in the UK. Not yet at least, despite it helping secure two patents in areas where its designer had no expertise at all, and was only responsible for feeding his algorithm the data. Legal arguments still deny AIs inventorship, since designers have responsibilities and are able to get into contracts, something an AI cannot do (Chen, 2020b). Additionally, AI might recognize patterns and generate predictions that were impossible before the age of big data. What is now being called "invisible data", such as spending behavior, or recognizing emotions from facial cues, are forms of data we never had access to before. In a way, AIs could generate data about consumers, to which the latter were never aware they had. In that case, ownership of that data might become murkier as AIs gain further acceptance as data creators, generators and maybe become entitled to some rights.

In that murky ethical space, many international figures from the tech scene to the catholic church have called on governments to intervene with regulation. These include Elon Musk calling for further scrutiny into his own company's AI development (Lyons, 2020), and the "Rome Call for AI Ethics" that includes six broad principles for ethical AI: transparency, inclusion, responsibility, impartiality, reliability, and security/privacy (Lyles, 2020). Journalists have naturally welcomed such statements but are also understandably unable to follow up with private tech companies as much as they would want to. The reason lies in the first and last principles mentioned above: Transparency versus privacy/security. In what is now known as the

"AI transparency paradox", companies need to maintain customers' data privately secure while being open about how they're doing it (Burt, 2019). This not only makes them vulnerable to security threats from hackers, but also from competition. For this reason, many AI products are protected as trade secrets and are therefore pretty much immune from public scrutiny. As an illustration, recently, amendments are being proposed to housing anti-discrimination laws in the US, where plaintiffs may need to prove 'intent to discriminate' to seek justice (Badger, 2019). As proprietary algorithms take over loan and housing risk score evaluations, it becomes much harder to prove and be compensated for housing discrimination because AIs inherently don't have intent. Journalists thereby need to embrace their role as an advocate for groups marginalized by AI systems, and as an adversary against technology companies that preach one thing and practice another.

4.3. AI as an ideology

A third issue that is unique to AI is the alienation of the citizens behind its development. As Weyl & Lanier put it, journalists ought not to "talk about how a machine is learning to see. Instead talk about how people contributed examples in order to define the visual qualities distinguishing "cats" from "dogs" in a rigorous way for the first time" (Lanier & Weyl, 2020). Unlike the few who are officially employed by the technology industry, those silent contributors are often neglected and therefore powerless in affecting change in the industry they ameliorate. The way various AI actors define work is warped, so that "when people provide data, behavioral examples, and even active problem solving online, it is not considered "work" but is instead treated as part of an off-the-books barter for certain free internet services"

(Lanier & Weyl, 2020). In that context, journalists have successfully highlighted the plight of social media platform moderators, since most of them are contractual and are often not paid proportionally to their role in their respective functions (Newton, 2019; Mack, 2019; Newton, 2020a). This was also brought up in the current 2020 COVID-19 pandemic, where social media platforms cautioned from potential mistakes in flagging content as they had to rely more on AI moderation given their reviewers could not get to their desks to review material that is too sensitive to be screened remotely (Kastrenakes, 2020). This still doesn't address the even larger labor behind the scenes generating big data in the first place. As mentioned earlier, as most AI stories are filed under the business section, much of the concerns and intricacies of decision making in this sector is hidden in the open from the general public's consideration. In that context, Brennen warns that industry indexing of most news stories about AI "undercut[s] consideration of the role of politics, public action, and collective decision-making in addressing AI" (Brennen, 2018). For this purpose, journalists can not only focus on industry workers in their coverage, but be more inclusive of regular citizens who are both customers and product contributors in this case.

AI discussions also suffer from overt technochauvinism, especially displayed when trying to regulate the field. As tech companies are understandably not enthusiastic about receiving blame for an AI that might not be working as intended, many are rushing to combat tech with more tech, while ignoring the underlying causes. For example, Facebook has recently moved to banning deepfakes, which are synthetic videos generated through AI. While deepfakes have been predominantly used in "creating" porn videos with female celebrities, very few have been

used for political purposes (Romano, 2019). It is obvious that all deepfakes are misleading, but not all misleading videos originate from AI. In fact, analysis on the podcast Reset describes this move from Facebook as a cop out, where they're banning software to avoid regulating speech (Duhaime Ross, 2020). As it is the case on most social media platforms, regulating misleading material on their sites not only requires bots and human supervision but is also a slippery slope. What qualifies as misleading is entirely up to the platform moderators' discretion, which warrants automatic censorship of offending content. However, public speech should be treated transparently, even if it can't be treated equally. The perils of such bans are even more pronounced in non-democracies. As China officially now criminalizes AI generated "fake news"(AFP, 2019), it becomes more urgent for journalists to carve out their roles as authenticators, as it becomes easier to criminalize dissent and controversial material based solely on its format and the whims of governments. For journalists covering AI, this could include distancing themselves from being mouthpieces and embracing the role of a missionary on the merits of transparency and free speech.

In other instances of technochauvinism, problems do exist in data, yet a tech fix seems to be easier than tackling underlying issues, especially in racial injustice cases. In the now infamous study, resumés sent out by people with black-sounding names had less chance to be selected for hiring than white-sounding counterparts. That finding was also replicated in a healthcare setting where a black diabetic patient received worse care than her white fellow patient. Sendhil Mullainathan, co-author on both studies, describes those two racial injustice incidents as similar with one big caveat: "In the first, hiring managers made biased decisions. In

the second, the culprit was a computer program, which recommended a less suitable medical action plan." She also admits that implicit bias training seems less effective than tweaking algorithms to make them less biased. She opines that "[c]hanging algorithms is easier than changing people: software on computers can be updated; the "wetware" in our brains has so far proven much less pliable." She foresees that with proper regulation, we might someday achieve less discriminatory algorithms than society ever was (Mullainathan, 2019). In fact, there's an industry that is selling AI generated faces, to boost the appearance of diversity in dating apps and HR promotional brochures among others (Harwell, 2020). Journalists therefore need to focus on how AI solutions work and more so on the shortcuts taken to fix them.

5. ALTERNATIVE FOCUSES: ANGLES THAT CAN BE APPLIED FROM TECH JOURNALISTS

5.1. Alternate Angle 1: Environmental costs

AI-driven curation certainly has its ethical limits, but also presents imminent material security threats that are not always highlighted by journalists. As a start, training AIs on expansive sets of data takes up computing power, which in some cases requires banks of computers running from days to weeks. In fact, a website was recently launched that lets AI researchers assess the projected carbon footprint of training their AI (Schmidt, Luccioni, Lacoste, & Dandres, 2020). Researchers at UMass Amherst found that training a single rather large AI model "may consume as much energy as a car over its entire lifetime—including the

energy needed to build it" (Knight, 2020a). This doesn't seem to be stopping as computing power required is increasing at a faster rate for each new AI landmark. As such, more articles highlighting how AI can help fight climate change need to inform readers of that limit (Snow, 2019). Concepts of technological somnambulism here suggest that journalists need not transmit a willingness to "sleepwalk through the process of reconstituting the conditions of human existence" (Langdon, 1986) and can act as an inspirator, highlighting consumer choices that are not guilt-ridden vis-à-vis our planet.

5.2. Alternate Angle 2: Concentration of power

Following that logic, the first physical limit may soon be dwarfed by the cost of running computer banks. A recent New York Times article highlights that it's becoming increasingly expensive to run those computers (Lohr, 2019). In turn, that might lead to a concentration of power with big companies, such as Google, Microsoft, Amazon and Facebook that can afford these data centers. This not only risks limiting competition to those who are already established, but also curbs AI research at universities. As it becomes more costly to train AIs that might not work, university research grants and their data centers simply cannot keep up. This will inevitably lead to fewer or less-trained researchers who will be someday joining the ranks of those companies (Lohr, 2019). Alternative evaluations of successful AIs are being proposed, such as Green AI, that aims to produce not only accurate but more efficient AI models in terms of training hours required (Schwartz, Dodge, Smith, & Etzioni, 2019). As this article is evidently categorized in the business section of the New York Times, such concerns ought to also be raised in subsequent discussions of the education sector and government

spending on AI research, questioning what it entails and who it enables. In fact, Brennen proposes that “collaborations between journalists on different beats could help outlets produce more sophisticated articles about AI” (Brennen, 2018). Even more ambitiously, Ball advocates for “scrap[ing] the idea of a “tech desk” altogether: The sector needs scrutiny, but since technology now touches every aspect of our society, keeping it siloed from the rest of the newsroom now feels artificial. Let it be covered, extensively, across desks” (Ball, 2018).

Even if journalists don’t emphasize concentration of clout in the world of AI, power holders will inevitably do it for them. In the biggest privacy scandal of 2020 so far, a startup called Clearview had been scraping data off of social media platforms, such as Facebook, Twitter and millions of websites to create a database of billions of identifiable faces (Hill, 2020). This example highlights the issue of AI as related to data privacy and governance. The scraping enabled Clearview to create a facial recognition app that was used by more than 600 law enforcement agencies and countless other firms and individuals. It was marketed as a crime solving tool, where users could start with a free-trial, upload images and usually get a match within minutes with a 75% accuracy. The story outraged many government officials, tech firms and concerned citizens, especially since the image bank is not government provided, nor was the security and accuracy of that startup’s technology ever externally scrutinized. The story keeps getting worse as the company was deciding who it was selling its tech to, and was refraining from selling it to “bad actors”. It is a private company’s right to be selective about its clientele, in the absence of regulation, but that can make it harder for journalists to trace to whom such powerful technology is being sold. As Clearview gained more exposure after that

article, it was successfully hacked weeks later (Yuen, 2020). Despite the hack only affecting their client list, this nonetheless is a reminder that concentration of power, in other words data, is a honey pot that might be irresistible to hackers. Yet, it seems that this fact was not obvious to law enforcement who willingly uploaded material to Clearview servers while using it on official searches. Even though the New York Times team was able to highlight this point, other reporters subsequently pressured their law enforcement agencies, including the RCMP, to not only admit to their use of the software but also to investigate any rights violations (Tunney, 2020).

During her investigation of Clearview, New York Times' Kashmir Hill asked some of her sources, law enforcement agents using the app, to query her picture. Most of the time this resulted in no results in the app, but at other times, it made her sources dodge her calls. As it turns out, Clearview had already internally flagged her as a person of interest, and would monitor users who might be searching for her. One user reported receiving a call from Clearview headquarters within minutes of his search to investigate why he'd been searching for this reporter. The reporter eventually confronted Clearview CEO of that finding, who smirked and apologized for that "software bug" while demonstrating live during the interview how it was now fixed and working (Brown, A., 2020). In other words, Clearview is watching all queries, manipulating results, and apparently flagging journalists who might be deemed troublesome.

Besides highlighting issues with surveillance and power dynamics, the Clearview story also highlights a common angle in tech stories: The inability to consent. In 2015, Alex Hern from the

Guardian undertook the monumental task of not ignoring end-user license agreements, the legal-term filled popup that most online users agree to without much thought, and read them (Hern, 2015). He did that experiment for a week, excluding all he'd had agreed to earlier in his online life, and ended up spending eight hours reading legal jargon. In short, he did that to highlight how unfeasible it is for regular netizens to process those agreements. Not only are these consent forms of contracts mostly in "legalese", but they're also lacking in granularity. Users either agree to the whole content, or decline and are denied access. In many cases surrounding AI, notably with Clearview, users may not even have the option to consent and are automatically entered into the system.

The inability to consent goes in tandem with the difficulty of opting out. For users whose data was scraped by Clearview, opting out is not guaranteed even if it's requested. In a recent Vox article on that topic, Heilweil indicates that opting out requests are only being processed if users reside in jurisdictions where "right to erasure" is a legal requirement (Heilweil, 2020). Moreover, "[i]n order to make those requests, you would need to confirm your identity with Clearview by sending in — get this — a photo of yourself on a government ID" (Heilweil, 2020). As such, journalists ought to use these occasions to emphasize the need for the legal framework surrounding such companies. For example, one journalist recently explored how easy it was to navigate the new California Consumer Privacy Act (CCPA) (Fowler, 2020). Theoretically, the CCPA should work similarly to the European General Data Protection Regulation (GDPR): broadly, users get to inquire what companies know about them, ask them not to sell their data, and even delete it. However, Fowler not only states the rights consumers

have but also attempts to negotiate with the tech companies. This process entailed her having to talk to each company separately to guarantee her rights not to have her data sold. The procedure also revealed that Google, Facebook and Amazon disagree that this portion of the act applies to them as they don't sell data. In fact, their business model is built on microtargeting users, which is not strictly selling the data, but selling the ability to manipulate users based on it. Exposing this loophole reminds readers that stricter legislation may be needed if these are the companies most citizens are worried about.

5.3 Alternate Angle 3: The need for granularity

That data privacy model may be tolerated by users so far, but is bound to raise more concern whenever more sensitive data is shared with third parties, or within the same system. For example, one should be able to allow full personal data sharing with an AI in a healthcare setting for diagnosis purposes, without extending those privileges to their billing department software. Besides the right to choose argument, two other reasons may motivate such a divergence. First, hackers may have fewer incentives to tap into medical software and cause misdiagnosis but more of them to tap into billing to funnel out some extra monetary compensation. Second, patients and other AI stakeholders do not have to enhance AI accuracy just because they can. This data-hungry endeavor needs to be examined not as an ultimate goal in all fields, but journalists ought to help foster discussions on which fields in each community would benefit from AI. Some societies may choose to steer away from developing smart bioweapons, while others would embrace it. Some societies may opt to invest more in AI security, to help protect their achievements from hacking while others opt to take more risks

for the sake of innovation. As Max Tegmark puts it, it's up to us to be complacent or ambitious about our goals for AI (Tegmark, 2018). In either case, it falls on good journalism to foster that kind of discussion.

5.4. Alternate Angle 4: Human stories

Maybe the most crucial of all in communicating AI to a general audience has to be its human element. In a haunting two-part podcast mini-series, Gabriel Dance and Michael Keller describe how they got tipped off about the world of "child sexual abuse imagery", avoiding the term porn as it might be confused with the legal adult industry and not the heinous crimes they are investigating (Barbaro, 2020). In this investigation, instead of calling on more efforts to detect this traffic online and ban it, the reporters explore how this might clash with another popular tech feature: encryption. Most consumers, especially journalists, go to certain lengths to secure their communications, hence the popularity of WhatsApp and Signal which offers end to end encryption. However, in cases of end to end encryption, abuse imagery can only be detected when it is shared via non-encrypted media (Facebook groups versus Facebook messenger). By exposing this dilemma with parents of a child sexual abuse victim, the reporters were able to re-focus listeners frenzy for privacy and encryption, and maybe make them reconsider how absolute they want it to be. Privacy and data security may be viewed as rights or products that could be sold, but in either case, fellow citizens need to be constantly informed as they debate this choice.

More urgently, citizens cannot afford to put this data rights conversation on their backburner. Google claims that it had the capacity to develop a Clearview-like software since 2011 but held back out of ethical reasons (Hill, 2020). Google had the upper hand and held back for a while, but could not have prevented Clearview from existing. Naturally, there's no monopoly on math and regulating AI is a race between technology and diplomacy. What makes this conversation more pressing, and probably more international in scale, is that large datasets are already available outside of North America. The secret ingredient to ameliorate AIs, big data, is already being harnessed in countries that are more populous, including India and China. The former is a democracy that has recently cut Kashmir off the internet for almost seven months, the longest in any democracy (Niha Masih, 2019), while the latter is a self-described "whole-process democracy"(Kim, 2019) . For this reason, advancing privacy concerns on how multi-national tech companies behave will not wait for major democracies for long and journalists ought to communicate that imperative.

5.5. Alternate Angle 5: Bottom-up approaches

Another angle that could be co-opted from tech would be how negotiating rights and their governance is rarely a one-size-fits-all deal. To illustrate this point, MIT launched an online quiz back in 2016 in which it gamified the classical trolley problem, giving players a choice of who to rescue from cyber death. The categories became more challenging with each stage, starting from a relatively straightforward kill one person or kill 3 to kill 2 elderly and save one child. 40 million ethical choices later from participants from more than 233 countries, their results were summed up in a 2018 paper which highlighted some universal preferences but also

cultural and economical-status specific variations (Vincent, J., 2018). Some were more careful about their elderly while others favored women over men. That was a strong ethical hint at how hard it is to implement a code of ethics, specifically in AI, that suits the worldwide population's beliefs. This same logic applies to AI governance, where Indians have recently proposed laws to protect their data from private companies, while being more lenient with their government's use of it (Goel, 2019). This is in stark contrast to the US, where some surveys indicate that citizens trust Amazon and Google more than their own government with their private data (Tiffany, 2018). This disparity between cultures also exists within larger countries, such as the US where the lack of a federal law concerning facial recognition has allowed states room to explore what their constituents want (Crawford, 2019). In those cases, journalists can champion their state's model if they think it's worthy, similarly to how healthcare and car emission standards were eventually nationalized.

5.6. Alternate Angle 6: Data journalism

Reporters have successfully leveraged data journalism to design interactive tools where readers could experience AI first hand. In this case, MIT TechReview published "Can you make AI fairer than a judge?" asking readers to go through a real-life example of COMPAS, a risk assessment algorithm commonly used in the US to evaluate risk score, and help judges make their final decision (Hao, 2020). Readers get to move sliders, draw the lines on threshold scores after which accused are incarcerated based on multiple definition of fairness. At the end, they are made to realize that keeping error rates comparable among groups (not too many are needlessly re-incarcerated) while treating people having the same risk score similarly is

impossible. The article cleverly then explains that this is the case not because those definitions of fairness are incompatible, but because of bias in existing historical data about black versus white people incarceration. It delivers this message not only by leveraging the power of data, but by presenting it as a form of actionable intelligence. Anyone can share this article/tool with a fellow citizen who is still unsure about the existence of this bias in the judicial system. Other attempts at addressing how the algorithm itself is not the problem can take a theoretical approach. In this case, Schneier opines that a ban on facial recognition technology doesn't work, as it only addresses one component of mass surveillance, identification, while ignoring correlation and discrimination (Schneier, 2020). Schneier argues that a broader discussion needs to happen with data brokers and regulators dealing with how private companies are allowed to discriminate. Journalism can function as a sense maker to help pinpoint how exactly a tech solution is not working, by leveraging data to illustrate it, and also how it could be remedied.

5.7. Alternate Angle 7: Keeping priorities straight even in emergencies

In light of recent social distancing measures for the 2020 COVID-19 pandemic, many have come in contact with technological questions they never had to grapple with earlier. For some laborers, this meant using software enabling work from home to keep themselves and the economy afloat. For suspected cases in isolation, this meant being tracked with government issued bracelets and accompanying apps. For many though, it also meant that our location data was pooled into "mobility reports" on countries and cities, to grade each on a sliding scale by how well they've managed to limit their movements (Google, 2020). However, many tech

journalists are skeptical of this tech powered approach to contain this health crisis. On the topic of data privacy, French parliamentarians have argued that tracking locations is not as helpful as it may seem, opining that what is needed is “un dépistage massif, et non d’un pistage massif dont nous aurons besoin” (Forteza, 2020). They explain how not only is the decision to track citizen location ethically dubious, as such decisions cannot be made in a time of crisis without parliamentary debate, but also practically limited. In their case, they think that it is an overreach to expect that enough people will download the app, as 60% of the population is needed for enough data gathering. This is especially complicated by the fact that the most vulnerable population, the elderly, is the least tech savvy. In that light, Casey Newton concedes there might be some merits to a “high-level heat map of human movements”, but also admits that the backlash on tech data collection practices, which have kept companies on the defensive for years, are currently on pause (Newton, 2020c).

This should not take away from the potential of AI in such crises, as it has successfully predicted which COVID-19 patients would develop serious respiratory diseases based on factors that were not obvious to practitioners (Devitt, 2020) . But, as with the previous example, it shouldn’t promote a blind trust in AI as the only way out of this situation. As Heaven puts it: “too much confidence in AI’s capabilities could lead to ill-informed decisions that funnel public money to unproven AI companies at the expense of proven interventions such as drug programs. It’s also bad for the field itself: overblown but disappointed expectations have led to a crash of interest in AI, and consequent loss of funding, more than once in the past” (Heaven, 2020). He acknowledges that available COVID-19 data pools, currently residing in different

hospital systems, public and private laboratories among others, are fertile ground for AI analysis but not at the moment. As highlighted earlier, big troves of data need to be available, streamlined and carefully checked for quality for a reliable model to emerge. Overall, these angles are just seven extant issues that point to strategies that are adopted in tech journalism in general, and could apply for journalists covering AI specifically.

6. PRACTICAL ADVICE FOR MOVING THE JOURNALISM CONVERSATION ON AI FORWARD

The successful AI journalism examples highlighted above can be built upon to deepen our discussion of data rights. These are seen here as particularly important due to their centrality in today's digital societies. As Seguar & Wasbord put it, it's in the space between data capitalism and data citizenship where sober discussions on the topic need to occur (Segura & Waisbord, 2019). By using the term "digital intelligence" to describe the synthesized data that AI builds from user data, Gurumurthy warns against their private ownership by corporations that generate them. She explains that "[w]hile private ownership of data per se may be legitimate, social advancement depends on a commons-based governance regime for social data" (Gurumurthy, 2018). This sentiment is also echoed by Montgomery who argues against private ownership of any sort as he reasons: "If the information is owned […] it will cease to belong to me provided that the transaction by which it is transferred is effective. Thus, it is inherent in the property paradigm that data 'about me' can 'cease to be 'mine' but this seems

to run against the expectations that those who assert that information is 'theirs' and provides a reason to be wary of the claim" (Montgomery, 2017).

As privacy and ownership rights only differ in their targets, information and possessions respectively, they are necessarily interdependent. As Kamleitner & Mitchell argue, "[t]he crossover between these rights becomes particularly pronounced in the case of information goods, such as software, books, or music" which in our context can also apply to big data (Kamleitner & Mitchell, 2019). These arguments have largely been underdeveloped by science journalists, as data ownership and privacy around AI systems is only brought up in the case of minors or deceased users (Barth, 2020; Brown, D. H. & Pecora, 2014). However, science journalists could borrow from parallel fields dealing with equally sensitive issues to drive this conversation forward.

6.1. The right to remain silent

The first example is to marry rights to privacy to the right to remain silent. This argument was brought up by Selinger & Hartzog in an article for the New York Times "What Happens When Employers Can Read Your Facial Expressions?"(Selinger & Hartzog, 2019). In this exploration, the authors argue against allowing the use of facial recognition in workplaces, as it might be used to reveal to employers not only your identity but also index specific facial cues to identify your emotions. This is clearly something most people are not willingly sharing, and could be expressed as the right to remain silent. In an even more progressive case, European countries have proposed principles for neurorights to regulate the use of brain data while Chile

is even advocating that the privacy of that data should be considered a human right (Samuel, 2019). As examples of commercial trading of brain data are already happening, advocacy groups saw the data's worth for marketing and its vulnerability for hacking, especially to provoke self-incrimination. For the above reasons, journalists can co-opt the right to privacy and to remain silent, and apply it to AI endeavors which might otherwise encroach on these important rights.

6.2. Learn from personalized medicine

The second example comes from personalized medicine, where ownership and privacy are still big issues. In that context, a lot of personalized medicine's success relies on genome sequencing. Once a genome is sequenced, treatment that is known to work well with that profile is prescribed. This matching can only be made more accurate as more people with that profile get their genome sequenced, which defines it as a data hungry endeavor (Anthony, 2019) . This poses inherent privacy risks not only to people's genetic data but to those who never consented to be tested, as one's DNA closely resembles their blood relatives. Forms of legislation, such as the US Genetic Information Nondiscrimination Act (Equal Employment Opportunity Commission, 2016) and Canadian Genetic Non-Discrimination Act (Standing Committee on Justice and Human Rights, 2017), exist to prevent discrimination on such basis but say little on genetic data ownership. As such, Montgomery argues that "[p]ublic ownership of genomic information is a more convincing model. It recognises obligations of stewardship to preserve and make productive use of the property without limiting the ability to secure a just distribution of the benefits" (Montgomery, 2017). Additionally, a recent US survey shows that

the genomic data conversation is now clearly framed as a right whose governance can be commercialized. Results from Penn State show that more than half of respondents were willing to provide their genomic data if they would gain monetary compensation while more than a third were unwilling under any circumstance (Briscoe, Ajunwa, Gaddis, & McCormick, 2020). Respondents were more willing to share if they had the right to be forgotten afterwards, and having reassurance that the data wouldn't be shared with a third party. Whether it's biometric or social data, delineating its possible misuse and ownership/privacy models doesn't require reinvention of the journalistic wheel.

6.3. Diversify coverage

As was highlighted earlier (see sections 4.1, 4.2, 5.2, 5.3), AI is a field that is high risk with many unknowns. Yet journalists communicating on the advances of AI still need to build a trusting relationship with their audience without resorting to overhyping, or dampening existing uncertainties. One way to appease skeptical audiences would be to clearly layout how AI could impact many aspects of their daily lives, and that broadness is reflected in the newsroom. In that line of thought, journalists can demonstrate how a datafied society is not only one that fights for data privacy and ownership but also interrogates how it can achieve data justice within it. The latter concept is an attempt by the Data Justice Lab to answer "how our understanding of social justice is changing in the context of datafication, what concepts and practices are needed, and how social justice can be advanced in a datafied society" (Dencik, Hintz, Redden, & Treré, 2019). This view acknowledges the different forces that shape this conversation, and can be applied in journalism by diversifying how AI articles are indexed. For

example, not all AI articles need to appear under the business section, and for that purpose, only target that audience. As more AI stories touch on a variety of societal issues, digital publications ought to adopt a more user-friendly search function for their articles. For example, readers could hypothetically search for “RCMP” and “Data Privacy” to get the article(s) that tackled both issues. On most platforms, searches are too wide and often lead to many unrelated articles. A more efficient curation and display of search results could signal to the reader that the newsroom is following up with said company or entity on specific issues. A similar tag could also be applied to journalists’ names when looking them up. In a system similar to ResearchGate, each investigator’s name has a portfolio of tags, where the nature and quantity depends on the work they’ve done: for example, journalists who cover Tesla and ethics often would have those among their top tags, helping readers establish trust in the fact that this beat is familiar to the reporter in question. All of these techniques are scalable, as once they’re implemented on an article level, they can help news consumers get a transparent overview and evaluate credibility accordingly.

6.4. Familiarize audiences with legislation and legal terminology

When privacy is being framed as a right, to be anonymous in public and to dissent, tech journalists need to familiarize their audience with legislation and legal terminology, and companies’ stances on them. For example, an article discussing Facebook in California should mention how the latter deals with CCPA requests to stop selling their data. Yet, not all rights need to be negotiated at the government level, some are even developing in urban areas. One of the pioneering ways by which the tech community is proposing to deal with data governance

is using data trusts. A data trust is how Sidewalk Labs plans to deal with urban data generated in its project with Waterfront Toronto. According to the Open Data Institute, a data trust is "a legal structure that provides independent third-party stewardship of data". Sidewalks defined theirs as a "civic data trust" considering it as a public asset, curtailing private ownership. In other words, data generated from this smart city will be protected, not by the Sidewalks, nor the government, but by a trust with its own custom rules. As this terminology seems to be catching on, there's no reason why tech journalists are not familiarizing the public with it. Unfortunately, this term is only mentioned a handful of times online, and only once in a story that wasn't about Sidewalk Labs (Austen, 2020; Vincent, D., 2019; Shee, 2019). Journalists covering AI may not even think to dive deeper into legislation and legal terminology, which returns us to the need to co-opt AI within a rights-based approaches to engage their audience. The essay ends with this co-option in mind.

7. CONCLUSION

By framing conversations about AI around privacy rights, journalists have a chance to engage and not just inform their audience. As many voices are now calling for a digital bill of rights to guarantee data privacy, journalists can also act as smart aggregators of how tech companies became so powerful in the first place (Srinivasan, 2020). In other words, journalists ought to help stakeholders avoid naïve realism, thinking that tech companies, citizens and governments are acting the way they do because they're uninformed, biased or irrational. As

highlighted by Casey Newton, there's a straight line from google maps to Clearview AI (Newton, 2020b). Newton rightly highlights that Google maps raised many eyebrows when it launched street view years ago. Unlike Clearview though, homeowners could always opt out and have their home removed from its directory. Yet, both technologies seized on publicly scrapable datasets and are in their own ways redefining the privacy conversation on a global scale. These depictions can serve as an additional call to resist our collective techno-somnambulism. Lanier & Weyl even argue that the nature of our conversation around AI should change. This technology is not only understood as a "number of competing philosophies that can direct our thinking about the nature and use of computation" but as a "political and social ideology rather than as a basket of algorithms" (Lanier & Weyl, 2020). Journalists are thereby empowered not only by ethical discussions in their academic circles, but wider privacy and ownership rights in legal and medical spheres. These domains should serve as different instruments that could be tapped into in order to reflect the diversity of the landscape in which AI news exist.

To summarize, the co-option of AI journalism within a rights-based approaches, with the central purpose of audience engagement, would include:

a) Guiding citizens in defining what problems are targeted by AI, and whether a tech solution or its scope are always suitable;
b) Avoiding AI overhype, as most advances in the field could imply an encroachment on either privacy or ownership rights;
c) Rallying the international tech journalist community to collaborate on covering multi-national AI endeavors, while also acknowledging that local solutions may be different;

d) Empowering netizens by publishing actionable intelligence outside of the business section;
e) Questioning technochauvinism, especially in light of environmental and democratic costs;
f) Highlighting the need for granularity not only in consent but also in balancing digital expectations such as encryption;
g) Exploring the limit of citizen governance within black-box AI companies, and how the latter's responsibilities could be evaded.

Despite the many warning signs, AI doomsday scenarios do not need to lead to self-fulfilling prophecies. Somehow refreshingly, recent research shows that bots amplify fake and real news equally (Vosoughi, Roy, & Aral, 2018), which puts the onus on journalists to be vigilant on what narratives and concerns deserve public attention in each AI story that is produced. AI is here to stay, but the way citizens advocate for their data rights within it remains flexible. In chapter 8 of her book, Broussard recommends that "[w]e should really focus on making human assistance systems instead of on making human-replacement systems. The point is not to make a world run by machines; people are the point. We need human-centered design" (Broussard et al., 2018). As everyone's privacy rights are potentially threatened by AI, journalism is uniquely positioned to leverage its storytelling and question if this reality is what society really wants.

REFERENCES

AFP. (2019, November 29). China criminalises 'fake news' created with AI, bots. *AlJazeera,* Retrieved from https://www.aljazeera.com/news/2019/11/china-criminalises-fake-news-created-ai-bots-191130062757933.html

Ali, W., & Hassoun, M. (2019). Artificial intelligence and automated journalism: Contemporary challenges and new opportunities. *Int J Media Journal Mass Commun, 5*(1), 40-49.

Anderson, J., & Rainie, L. (2020). *Many tech experts say digital disruption will hurt democracy.* ().Pew Research Center. Retrieved from https://www.pewresearch.org/internet/2020/02/21/many-tech-experts-say-digital-disruption-will-hurt-democracy/

Anthony, A. (2019, September 28). Will genome sequencing bring precision medicine for all? *The Guardian,* Retrieved from https://www.theguardian.com/science/2019/sep/28/genome-sequencing-precision-medicine-bespoke-healthcare-nhs

Austen, I. (2020, March 10). You can't fight city hall. but maybe you can fight google. *The New York Times* Retrieved from https://www.nytimes.com/2020/03/10/world/canada/toronto-sidewalk-labs-google.html

Badger, E. (2019, August 20). Who's to blame when algorithms discriminate? *The New York Times* Retrieved from https://www.nytimes.com/2019/08/20/upshot/housing-discrimination-algorithms-hud.html

Ball, J. (2018). We need a new model for tech journalism. *Columbia Journalism Review,*

The Daily (Producer), & Barbaro, M. (Director). (2020, February 19). *A criminal underworld of child abuse.* [Video/DVD] Retrieved from https://www.nytimes.com/2020/02/19/podcasts/the-daily/child-sex-abuse.html

Barth, B. (2020, April 08). The digital afterlife. *The Walrus,* Retrieved from https://thewalrus.ca/the-digital-afterlife/

Biswal, S. K., & Gouda, N. K. (2020). Artificial intelligence in journalism: A boon or bane? *Optimization in machine learning and applications* (pp. 155-167) Springer.

Brennen, J. (2018). An industry-led debate: How UK media cover artificial intelligence.

Briscoe, F., Ajunwa, I., Gaddis, A., & McCormick, J. (2020). Evolving public views on the value of oneâ€™s DNA and expectations for genomic database governance: Results from a national survey. *PloS One, 15*(3), e0229044.

Broussard, M., Diakopoulos, N., Guzman, A. L., Abebe, R., Dupagne, M., & Chuan, C. (2019). Artificial intelligence and journalism. *Journalism & Mass Communication Quarterly, 96*(3), 673-695.

Broussard, M., IEEE Xplore, & MIT Press. (2018). *Artificial unintelligence :How computers misunderstand the world*. Cambridge, Massachusetts; Piscataway, New Jersey: The MIT Press; IEEE Xplore.

The Daily (Producer), & Brown, A. (Director). (2020, February 10). *The end of privacy as we know it?* [Video/DVD] Retrieved from https://www.nytimes.com/2020/02/10/podcasts/the-daily/facial-recognition-surveillance.html

Brown, D. H., & Pecora, N. (2014). Online data privacy as a children's media right: Toward global policy principles. *Journal of Children and Media, 8*(2), 201-207.

Burt, A. (2019, December 13). The AI transparency paradox. *Harvard Business Review,* Retrieved from https://hbr.org/2019/12/the-ai-transparency-paradox

Campbell, A. F. (2018, October 18). How tech employees are pushing silicon valley to put ethics before profit. *Vox,* Retrieved from https://www.vox.com/technology/2018/10/18/17989482/google-amazon-employee-ethics-contracts

Chen, A. (2018, May 23). How computers misunderstand the world. *The Verge,* Retrieved from https://www.theverge.com/2018/5/23/17384324/meredith-broussard-artifical-unintelligence-technology-criticism-technochauvinism

Chen, A. (2020a, January 28). 40 groups have called for a US moratorium on facial recognition technology. *MIT Technology Review,* Retrieved from https://www.technologyreview.com/f/615098/facial-recognition-clearview-ai-epic-privacy-moratorium-surveillance/

Chen, A. (2020b, January 09). Can an AI be an inventor? not yet. *MIT Technology Review,* Retrieved from https://www.technologyreview.com/s/615020/ai-inventor-patent-dabus-intellectual-property-uk-european-patent-office-law/

Chuan, C., Tsai, W. S., & Cho, S. Y. (2019). Framing artificial intelligence in american newspapers. Paper presented at the *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society,* 339-344.

Chyi, H. I., & Lee, A. M. (2018). Commercialization of technology news: Quantity and quality dimensions. *Journalism Practice, 12*(5), 585-604.

Crawford, S. (2019, December 13). Facial recognition laws are (literally) all over the map. *Wired,* Retrieved from https://www.wired.com/story/facial-recognition-laws-are-literally-all-over-the-map/

Dencik, L., Hintz, A., Redden, J., & Treré, E. (2019). Exploring data justice: Conceptions, applications and directions. *Information, Communication & Society, 22*(7), 873-881. doi:10.1080/1369118X.2019.1606268

Devitt, J. (2020, March 30). Tool predicts who'll get seriously sick from COVID-19. *Futurity,* Retrieved from https://www.futurity.org/covid-19-artificial-intelligence-severe-respiratory-disease-2320402/

Reset (Producer), & Duhaime Ross, A. (Director). (2020, January 16). *Facebook's fake deepfake problem.* [Video/DVD] Vox. Retrieved from https://player.fm/series/reset-2555404/facebooks-fake-deepfake-problem

Genetic information nondiscrimination act, (2016). Retrieved from https://www.federalregister.gov/documents/2016/05/17/2016-11557/genetic-information-nondiscrimination-act

Fairfield, J., & Shtein, H. (2014). Big data, big problems: Emerging issues in the ethics of data science and journalism. *Journal of Mass Media Ethics, 29*(1), 38-51. doi:10.1080/08900523.2014.863126

Forteza, P. (2020, April 09). Tribune de 15 députés: "Nous avons besoin d'un dépistage massif, pas d'un pistage massif". *Le Figaro,* Retrieved from https://www.lefigaro.fr/vox/politique/tribune-de-15-deputes-nous-avons-besoin-d-un-depistage-massif-pas-d-un-pistage-massif-20200408

Fowler, G. A. (2020, February 19). Don't sell my data! we finally have a law for that. *The Washington Post* Retrieved from https://www.washingtonpost.com/technology/2020/02/06/ccpa-faq/?arc404=true

Friedman, S. M., Dunwoody, S., & Rogers, C. L. (1999). *Communicating uncertainty :Media coverage of new and controversial science*. Mahwah, N.J.: L. Erlbaum Associates.

Garvey, C., & Maskal, C. (2019). Sentiment analysis of the news media on artificial intelligence does not support claims of negative bias against artificial intelligence. *Omics: A Journal of Integrative Biology,*

Goel, V. (2019, December 10). On data privacy, india charts its own path. *The New York Times* Retrieved from https://www.nytimes.com/2019/12/10/technology/on-data-privacy-india-charts-its-own-path.html

Google. (2020). COVID-19 community mobility report. Retrieved from https://www.google.com/covid19/mobility/

Gurumurthy, A. (2018). Where is the 'struggle'in communications for social progress? *Global Media and Communication, 14*(2), 193-200.

Gutierrez-Lopez, M., Missaoui, S., Makri, S., Porlezza, C., Cooper, G., & MacFarlane, A. (2019). Journalists as design partners for AI. Paper presented at the *Workshop for Accurate, Impartial and Transparent Journalism: Challenges and Solutions. CHI 2019.*

Haider, M., & Gandomi, A. (2019). When big data made the headlines: Mining the text of big data coverage in the news media. *International Journal of Services Technology and Management, Forthcoming,*

Hanitzsch, T., & Vos, T. P. (2018). Journalism beyond democracy: A new look into journalistic roles in political and everyday life. *Journalism, 19*(2), 146-164.

Hansen, M., Roca-Sales, M., Keegan, J. M., & King, G. (2017). Artificial intelligence: Practice and implications for journalism.

Hao, K. (2019, August 03). China has started a grand experiment in AI education. it could reshape how the world learns. *MIT Technology Review,* Retrieved from https://www.technologyreview.com/s/614057/china-squirrel-has-started-a-grand-experiment-in-ai-education-it-could-reshape-how-the/

Hao, K. (2020, January 10). Can you make AI fairer than a judge? play our courtroom algorithm game. *MIT Technology Review,* Retrieved from https://www.technologyreview.com/s/613508/ai-fairer-than-judge-criminal-risk-assessment-algorithm/

Harwell, D. (2020, January 07). Dating apps need women. advertisers need diversity. AI companies offer a solution: Fake people. *The Washington Post* Retrieved from https://www.washingtonpost.com/technology/2020/01/07/dating-apps-need-women-advertisers-need-diversity-ai-companies-offer-solution-fake-people/

Heaven, W. (2020, March 12). AI could help with the next pandemic-but not with this one. *MIT Technology Review,* Retrieved from https://www.technologyreview.com/s/615351/ai-could-help-with-the-next-pandemicbut-not-with-this-one/

Heilweil, R. (2020, March 06). The world's scariest facial recognition app company keeps lying. *Vox,* Retrieved from https://www.vox.com/recode/2020/2/11/21131991/clearview-ai-facial-recognition-database-law-enforcement

Hern, A. (2015, June 15). I read all the small print on the internet and it made me want to die. *The Guardian,* Retrieved from https://www.theguardian.com/technology/2015/jun/15/i-read-all-the-small-print-on-the-internet

Hill, K. (2020, January 18). The secretive company that might end privacy as we know it. *The New York Times* Retrieved from https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html

IEEE position statement artificial intelligence (2019). Retrieved from https://globalpolicy.ieee.org/wp-content/uploads/2019/06/IEEE18029.pdf

Kamleitner, B., & Mitchell, V. (2019). Your data is my data: A framework for addressing interdependent privacy infringements. *Journal of Public Policy & Marketing, 38*(4), 433-450.

Kastrenakes, J. (2020, March 16). YouTube will rely more on AI moderation while human reviewers can't come to the office. *The Verge,* Retrieved from https://www.theverge.com/2020/3/16/21182011/youtube-ai-moderation-coronavirus-video-removal-increase-warning

Katyal, S. K. (2019). Private accountability in the age of artificial intelligence. *UCLA L.Rev., 66*, 54.

Kim, J. (2019, December 06). Exploring china's new narrative on democracy. *The Diplomat,* Retrieved from https://thediplomat.com/2019/12/exploring-chinas-new-narrative-on-democracy/

Knight, W. (2020a, January 21). AI can do great things-if it doesn't burn the planet. *Wired,* Retrieved from https://www.wired.com/story/ai-great-things-burn-planet/

Knight, W. (2020b, February 22). This technique uses AI to fool other AIs. *Wired,* Retrieved from https://www.wired.com/story/technique-uses-ai-fool-other-ais/

Langdon, W. (1986). The whale and the reactor: A search for limits in an age of high technology.

Lanier, J., & Weyl, G. (2020, March 15). AI is an ideology, not A technology. *Wired,* Retrieved from https://www.wired.com/story/opinion-ai-is-an-ideology-not-a-technology/

Lewis, S. C., & Westlund, O. (2015). Big data and journalism: Epistemology, expertise, economics, and ethics. *Digital Journalism, 3*(3), 447-466.

Liu, S. (2020). Global AI funding by category 2019. Retrieved from https://www.statista.com/statistics/943136/ai-funding-worldwide-by-category/

Lohr, S. (2019, September 26). At tech's leading edge, worry about a concentration of power. *The New York Times* Retrieved from https://www.nytimes.com/2019/09/26/technology/ai-computer-expense.html

Lyles, T. (2020, February 28). The catholic church proposes AI regulations that "protect people". *The Verge,* Retrieved from http://www.theverge.com/2020/2/28/21157667/catholic-church-ai-regulations-protect-people-ibm-microsoft-sign

Lyons, K. (2020, February 18). Elon musk says AI development should be better regulated, even at tesla. *The Verge,* Retrieved from https://www.theverge.com/2020/2/18/21142489/elon-musk-ai-regulation-tweets-open-ai-tesla-spacex-twitter

Mack, Z. (2019, July 02). Why AI can't fix content moderation. *The Verge,* Retrieved from https://www.theverge.com/2019/7/2/20679102/content-moderation-ai-social-media-behind-the-screen-sarah-t-roberts-vergecast

Marr, B. (2018, February 14). The key definitions of artificial intelligence (AI) that explain its importance. *Forbes,* Retrieved from https://www.forbes.com/sites/bernardmarr/2018/02/14/the-key-definitions-of-artificial-intelligence-ai-that-explain-its-importance/#3559f8984f5d

Menand, L. (2018, June 11). Why do we care so much about privacy? *The New Yorker,* Retrieved from https://www.newyorker.com/magazine/2018/06/18/why-do-we-care-so-much-about-privacy

Molnar, P. (2019). Technology on the margins: AI and global migration management from a human rights perspective. *Cambridge International Law Journal, 8*(2), 305-330.

Montgomery, J. (2017). Data sharing and the idea of ownership. *The New Bioethics, 23*(1), 81-86.

Mullainathan, S. (2019, December 06). Biased algorithms are easier to fix than biased people. *The New York Times* Retrieved from https://www.nytimes.com/2019/12/06/business/algorithm-bias-fix.html

Nechushtai, E., & Lewis, S. C. (2019). What kind of news gatekeepers do we want machines to be? filter bubbles, fragmentation, and the normative dimensions of algorithmic recommendations. *Computers in Human Behavior, 90*, 298-307.

Newton, C. (2019, February 25). The secret lives of facebook moderators in america. *The Verge,* Retrieved from https://www.theverge.com/2019/2/25/18229714/cognizant-facebook-content-moderator-interviews-trauma-working-conditions-arizona

Newton, C. (2020a, March 20). Facebook is nearing a settlement with its content moderators in a class action lawsuit. *The Verge,* Retrieved from https://www.theverge.com/2020/3/20/21188356/facebook-content-moderator-lawsuit-class-action-settlement

Newton, C. (2020b, February 07). The straight line from google maps to clearview AI. *The Verge,* Retrieved from https://www.theverge.com/interface/2020/2/7/21126498/google-maps-birthday-privacy-street-view-germany-clearview-ai

Newton, C. (2020c, April 07). Tech giants are finding creative ways to use our data to fight the coronavirus. *The Verge,* Retrieved from https://www.theverge.com/interface/2020/4/7/21209595/google-mobility-reports-facebook-disease-prevention-maps-tech-privacy-backlash

Niha Masih, S. I. (2019, December 16). India's internet shutdown in kashmir is the longest ever in a democracy. *The Washington Post* Retrieved from https://www.washingtonpost.com/world/asia_pacific/indias-internet-shutdown-in-kashmir-is-now-the-longest-ever-in-a-democracy/2019/12/15/bb0693ea-1dfc-11ea-977a-15a6710ed6da_story.html

O'Brien, D. (2020, January 29). Data privacy or data protection day? it's a human right, either way. *The Electronic Frontier Foundation,* Retrieved from https://www.eff.org/deeplinks/2020/01/data-privacy-or-data-protection-day-its-human-right-either-way

Raval, T. (2019, November 12). Council post: Data privacy as A basic human right. *Forbes,* Retrieved from https://www.forbes.com/sites/forbestechcouncil/2019/11/12/data-privacy-as-a-basic-human-right/#5c1e01a64cbf

Romano, A. (2019, October 07). Deepfakes are a real political threat. for now, though, they're mainly used to degrade women. *Vox,* Retrieved from https://www.vox.com/2019/10/7/20902215/deepfakes-usage-youtube-2019-deeptrace-research-report

Russell, S. (2019, October 08). How to stop superhuman A.I. before it stops us. *The New York Times* Retrieved from https://www.nytimes.com/2019/10/08/opinion/artificial-intelligence.html

Samuel, S. (2019, December 20). Brain-reading tech is coming. the law is not ready to protect us. *Vox,* Retrieved from https://www.vox.com/2019/8/30/20835137/facebook-zuckerberg-elon-musk-brain-mind-reading-neuroethics

Satariano, A. (2019, September 15). Real-time surveillance will test the british tolerance for cameras. *The New York Times* Retrieved from https://www.nytimes.com/2019/09/15/technology/britain-surveillance-privacy.html

Schmidt, V., Luccioni, A., Lacoste, A. & Dandres, T. (2020). Machine learning CO2 impact calculator. Retrieved from https://mlco2.github.io/impact/#home

Schneier, B. (2020, January 20). We're banning facial recognition. we're missing the point. *The New York Times* Retrieved from https://www.nytimes.com/2020/01/20/opinion/facial-recognition-ban-privacy.html

Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2019). Green AI. *ArXiv.Org,*

Segura, M. S., & Waisbord, S. (2019). Between data capitalism and data citizenship. *Television & New Media, 20*(4), 412-419.

Selinger, E., & Hartzog, W. (2019, October 17). What happens when employers can read your facial expressions? *The New York Times* Retrieved from https://www.nytimes.com/2019/10/17/opinion/facial-recognition-ban.html

The AI Element (Producer), & Shee, A. (Director). (2019, September 19). *In data we trust.* [Video/DVD] Retrieved from https://www.elementai.com/podcast#taking-control-of-our-data

Smith, B., & Browne, C. A. (2019, September 07). Do our faces deserve the same protection as our phones? *Wired,* Retrieved from https://www.wired.com/story/do-our-faces-deserve-the-same-protection-as-our-phones/

Snow, J. (2019, July 18). How artificial intelligence can tackle climate change. *National Geographic,* Retrieved from https://www.nationalgeographic.com/environment/2019/07/artificial-intelligence-climate-change/#close

Srinivasan, R. (2020, January 28). Americans need a 'digital bill of rights'. here's why. *The Guardian,* Retrieved from https://www.theguardian.com/commentisfree/2020/jan/28/americans-need-a-digital-bill-of-rights-heres-why#maincontent

Genetic non-discrimination act, (2017). Retrieved from https://laws-lois.justice.gc.ca/eng/acts/G-2.5/index.html

Stray, J. (2019). Making artificial intelligence work for investigative journalism. *Digital Journalism, 7*(8), 1076-1097.

TED (Producer), & Tegmark, M. (Director). (2018, July 05). *How to get empowered, not overpowered, by AI.* [Video/DVD] Retrieved from https://www.youtube.com/watch?v=2LRwvU6gEbA

Therrien, D. (2020). A data privacy day conversation with canada's privacy commissioner. Retrieved from https://www.priv.gc.ca/en/opc-news/speeches/2020/sp-d_20200128/

Tiffany, K. (2018, October 25). America is losing faith in its institutions - but not in amazon. *Vox,* Retrieved from https://www.vox.com/the-goods/2018/10/25/18022956/amazon-trust-survey-american-institutions-ranked-georgetown

Tunney, C. (2020, March 05). RCMP denied using facial recognition technology - then said it had been using it for months | CBC news. *Cbc,* Retrieved from https://www.cbc.ca/news/politics/clearview-ai-rcmp-facial-recognition-1.5482266

Vergeer, M. (2020). Artificial intelligence in the dutch press: An analysis of topics and trends. *Communication Studies,* , 1-20.

Vincent, D. (2019, September 26). Sidewalk labs' urban data trust is 'problematic,' says ontario privacy commissioner. *The Star* Retrieved from

https://www.thestar.com/news/gta/2019/09/26/sidewalk-labs-urban-data-trust-is-problematic-says-ontario-privacy-commissioner.html

Vincent, J. (2018, October 24). Global preferences for who to save in self-driving car crashes revealed. *The Verge,* Retrieved from https://www.theverge.com/2018/10/24/18013392/self-driving-car-ethics-dilemma-mit-study-moral-machine-results

Vincent, J. (2019, January 28). The state of AI in 2019. *The Verge,* Retrieved from https://www.theverge.com/2019/1/28/18197520/ai-artificial-intelligence-machine-learning-computational-science

Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science (New York, N.Y.), 359*(6380), 1146-1151. doi:10.1126/science.aap9559 [doi]

Yuen, J. (2020, February 28). Entire client list hacked in clearview AI data breach. *Toronto Sun,* Retrieved from https://torontosun.com/news/local-news/entire-client-list-hacked-in-clearview-ai-data-breach