



An Industry-Led Debate: How UK Media Cover Artificial Intelligence

Authors: J. Scott Brennen, Philip N. Howard, and Rasmus Kleis Nielsen

Key findings

As industry, government, and academia invest in various forms of artificial intelligence, many believe that these rapidly developing technologies will radically reshape life across the United Kingdom. How they might do so, however, and what role the public might play in shaping that transformation, all remain open questions. Quality news coverage is essential to the vibrant and critical public discussion needed to confront this emerging public issue. In this RISJ Factsheet, we analyse eight months of reporting on artificial intelligence (AI) in six mainstream news outlets in the United Kingdom. Our mixed-methods analysis of 760 articles that reference AI reveals three main findings:

- Nearly 60 percent of news articles across outlets are indexed to industry products, initiatives, or announcements. 33 percent of unique sources across all articles are affiliated with industry, almost twice as many as those from academia, and six times as many as those from government. Nearly 12 percent of all articles reference Elon Musk.
- Portraying AI as a relevant and competent solution to a range of public problems, outlets regularly assert the influence it will have across areas of public life often with little acknowledgement of on-going debates concerning AI's potential effects.
- As an emerging public issue, AI is being politicised through the topics that outlets emphasise in their coverage:

- Right-leaning outlets highlight issues of economics and geopolitics, including automation, national security, and investment.
- Left-leaning outlets highlight issues of ethics of AI, including discrimination, algorithmic bias, and privacy.

While news coverage provides an important foundation for public discussion of AI, experts continue to disagree about what AI is, what it will be able to do, and how it can be designed, regulated, and integrated into society. The recognition by news outlets that there are legitimately different political interpretations of what AI is and what topics deserve public attention might help bridge the diverse conversations occurring around AI and facilitate a richer public dialogue. Public discussion might also benefit from news outlets moving beyond industry initiatives and sources that tend to focus on only one side and thus can undercut a wider understanding of AI as a public issue. Increased engagement from scientists, activists, and others can provide alternative and independent views on the capabilities, promises, and pitfalls of AI, while helping portray AI less as a world-ending disruption and more as a set of powerful new technologies that are in the process of being developed.

General Overview

Artificial intelligence is a term that is both widely used and loosely defined. Most basically, AI is a collection of ideas, technologies, and techniques that relate to

a computer system's capacity to, as Dickens Olewe of the BBC described it, 'perform tasks normally requiring human intelligence' (5 Apr. 2018). Research on artificial intelligence dates to the beginning of computing. However, over the last 15 years, there have been major advancements in the field as a result of increased computational processing power, developments in algorithms, and perhaps most importantly, the availability of large data sets that can help train AI systems (Select Committee on AI, 2018). Today, most AI systems involve machine or deep learning, types of algorithms that can both recognise patterns in data sets with little human direction and improve over time.

Companies across the world are investing heavily in artificial intelligence research and development. Both new start-ups and existing companies are integrating artificial intelligence into a wide range of products from self-driving cars, to weapons, to athletic shoes – and services, including health care, news production, and social media content curation. Over the last few years, the UK government has also made artificial intelligence a major policy initiative. In May 2018, the government announced the Artificial Intelligence and Data Grand Challenge, an effort 'to put the UK at the forefront of the AI and data revolution' (Gov.uk, 2018). In recent speeches, Theresa May has identified AI as a major growth area for both British industry and health care. That being said, there remains significant uncertainty and debate concerning not only what effects AI may have across society (Select Committee on AI, 2018), but also how we should regulate and develop AI systems to ensure their equitable and safe deployment for the public good (see Dafoe, 2018).

Amid notable industry, government, and academic interest, AI has become a popular topic in UK news. However, there have been few systematic analyses of how UK media are covering AI. A House of Lords report released in April 2018 included testimony from several journalists and AI experts on media coverage. Opinions ranged from that of BBC technology reporter Rory Cellan-Jones, who observed, 'I think we are doing a pretty broad and, generally, sensible job, with the occasional bout of alarmism' (Select Committee on AI, 2017, Q.13), to the computer scientist Peter McOwan, who argued that public discussion is

overweighed at the moment by the negative stories associated with AI. Very often these negative stories are somewhat sensational, not surprisingly because they are picked up by the newspapers and very often

are not based on the credible technicalities of what is available to us at the moment and include a very large pinch of future gazing. (Q.214)

Others have argued that media coverage frequently swings between two sensational poles: utopian dreams of workless futures and eternal life, and dystopian nightmares of robot uprisings and the apocalypse (Craig, 2018).

Mainstream media coverage of AI is developing against a backdrop of structural changes across the news industry, including persistent economic disruption and the digital transformation (Kuong, 2017; Newman et al., 2018). Specialty reporting – including science and technology journalism – has been especially impacted. Some outlets have reduced or even eliminated their science and/or technology desks. These changes mean that some outlets cover these stories less frequently, task non-specialist reporters with reporting these stories, give their reporters less time and fewer resources to cover them, or encourage more reliance on press releases or wire articles (Schäfer, 2017; Dunwoody, 2014). These pressures and challenges all complicate reporting on a topic as new and technically complex as AI.

Despite these myriad challenges, mainstream news outlets remain a key space for, and influence on, public discussion. As AI spreads into diverse areas of public life through new products, major research initiatives, and automated decision-making, we need to understand better how technical research and expert views are translated into public. But we also need to understand better who is being given space to discuss AI and what they are saying: the public narratives, expectations, hopes, and fears surrounding AI. News coverage can provide publics with space and resources to make sense of and address pressing public problems. Studying media discussion of AI helps elucidate what AI is, what AI could be, and what AI means to publics.

The findings described here derive from a systematic analysis of a corpus of 760 articles produced in the first eight months of 2018 by six mainstream UK news outlets. The corpus comprises all written content produced and archived by these outlets, including news, features, and commentary. These outlets were strategically selected to represent a variety of political leanings as well as a mix of legacy and digital-born outlets. The corpus includes two right-leaning outlets, the *Telegraph*, and the *MailOnline*¹; two left-leaning

¹ Given that the *MailOnline* and the *Daily Mail* have distinct editorial teams and approaches, we look at *both* outlets separately.

outlets, the *Guardian*, and the HuffPost; one publicly funded outlet, the BBC; and one technology-specific outlet, the UK edition of Wired. Articles were collected either from the LexisNexis archive or from outlets' own online archives through targeted searches of the phrases: 'artificial intelligence', 'machine learning', 'deep learning', and 'neural networks'.² Articles were coded for a range of data including outlet, author(s), article type (news, feature, opinion, etc.), and news peg (academic study, product release, government report, etc.). Also, every unique direct quote within articles was coded by type of source (e.g. academic computer scientist, politician, CEO, etc.). Finally, articles were inductively coded for recurring themes, topics, and frames.

Three Themes

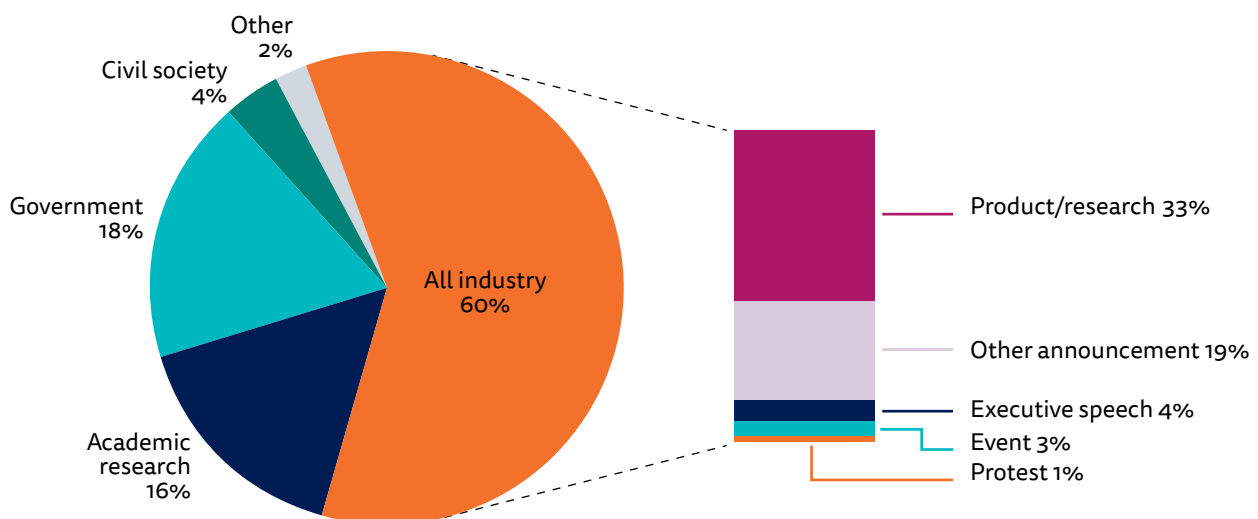
Our analysis reveals a number of distinct themes within the corpus, ranging from AI's potential in healthcare to automation to global competition in AI development. Rather than describe each in turn, we select three of the most common: one that is found across outlets, one that is more common in right-leaning outlets, and one that is more common in left-leaning outlets. Together, these three themes highlight not only the incipient politicisation of AI as a public issue, but also the predominance of frames that prioritise industry initiatives while positioning AI as a widely relevant and competent solution to a variety of public problems.

New Industry Products, Announcements, and Research

By a notable margin, the single most common topic or theme across coverage involves introducing, reviewing, or critiquing commercial products, initiatives, or research. Nearly 60 percent of news articles were coded as being framed around an industry product. As news pegs, new industry products or initiatives far outpace academic studies, reports, or political speeches (see Figure 1).

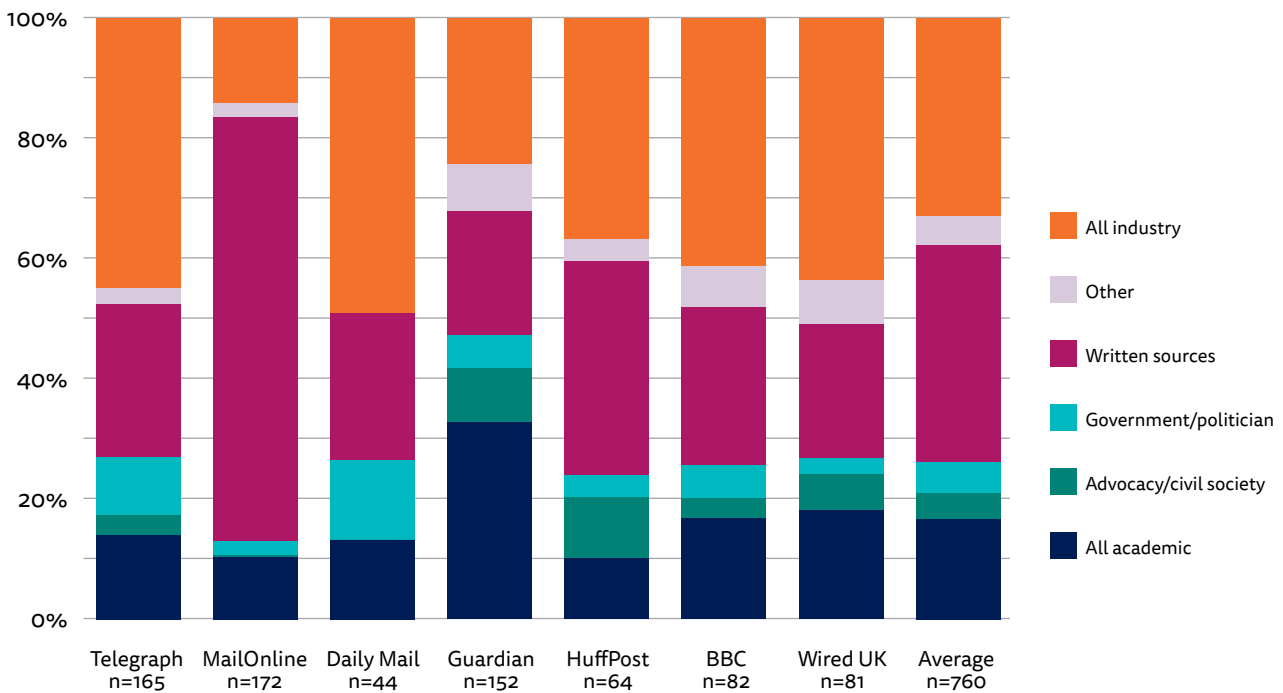
In being indexed to industry concerns, many news and feature stories describe new products that include artificial intelligence. Products range from those as mundane as smart phones, or running shoes, to those as outrageous as sex robots or brain preservation. Many articles follow business dealings or AI-related initiatives of large technology companies. Start-ups, buyouts, and investments all generated coverage, as did ongoing efforts such as Facebook's AI-driven content moderation or Google's DeepMind. Outlets also regularly covered industry promotional events, such as IBM's 'debate' between humans and an AI, or conferences and tech shows like the Consumer Electronics Show (CES) or the developer conference Google I/O. High-level executives at tech companies involved in AI also regularly inspired coverage. Whether this was Elon Musk pontificating about the future of AI, or when Jane Wakefield of the BBC reported, 'DeepMind co-founder Shane Legg gives

Figure 1: Relative proportion of news pegs in news articles



(n=419). 'Academic research' news pegs include speeches by academic researchers and release of findings from academic studies. 'Government' includes political speeches and government reports. 'All industry' includes products, announcements, business dealings, and research.

² That being said, nearly every article in the corpus includes the phrase 'artificial intelligence.' Even those articles that discuss machine learning, deep learning, or neural networks, specifically identify these as techniques of AI.

Figure 2: Relative percentages of unique sources in all articles across outlets

'All industry' includes company executives, researchers, employees, and spokespersons. 'Written sources' include quotations from press releases, official statements, speeches, as well as from other news articles. 'Government/politician' includes politicians, civil servants, MPs, and ministers. 'Advocacy/civil society' includes members of advocacy organisations. 'All academic' includes any researcher or administrator employed by a university.

London teen top AI tips' (2 Aug. 2018). Finally, just over 2 percent of the corpus, or 17 articles across four outlets, were identified as being paid for and placed by companies. Sponsoring companies included Barclays, Pharmacy2U, J.P. Morgan, Hitachi, and Sandoz and non-profit organisations such as the Thomson Reuters Foundation, Rockefeller, and Skoll Foundations.

Industry influence over AI coverage is not, however, limited to story topics: industry-connected sources also predominate across outlets. As seen in Figure 2, 33 percent of unique sources across outlets are industry related. This is almost twice the proportion of academic sources and six times more than political or government sources.³ The vast majority of these industry sources are CEOs or other high-level executives.⁴ Notably, the *Guardian* is the only outlet in which industry sources were not the most cited; academic sources were the most cited in the *Guardian* (32.7 percent).

Several persistent frames emerged across the hundreds of articles detailing industry concerns. First, across news, feature, and opinion articles, products

are often framed as solutions to on-going problems. These problems range from cancer to renewable energy to road rage to judging 'If Your Outfit is Good or Bad' (HuffPost UK, 7 June 2018) or keeping 'the passion in a relationship' (HuffPost UK, 10 Apr. 2018). Taken together, the implication becomes that all these different types of problems are best approached not only through a technological solution but through an AI-driven technological solution. Rarely do journalists or commentators question if (AI-containing) new technologies are the *best* solutions to these myriad problems.

Second, AI is frequently recognised as bringing massive changes across sectors, from revolutionising the mining industry and warfare to transforming fish farming and healthcare. The MailOnline is especially fond of heralding, and quoting those that herald, AI-led revolutions in everything from 'how the world's music is organized and curated' (23 May, 2018) to 'man's relationship with technology' (9 Jan. 2018). Some outlets amplify the potential implications of these products by focusing on either the *intention* behind or the *potential* of a new product, rather than its

³ This broad category includes politicians, civil servants, MPs, and government ministers.

⁴ The relative source makeup varied across outlets. The highest category overall was coded as 'written sources', this included quotations from press releases, speeches, statements, or other news stories. This was notably high for the MailOnline (70.5 percent).

current functionality. When the *Telegraph* notes that ‘Viagra inventor aims to heal “broken” drugs model with AI’ (26 July 2018), or that ‘Google is planning to use satellite imagery to map the “solar potential” of Britain’s rooftops’ (31 Mar. 2018) it obscures the distinction between what is actually possible and what is aspirational. Perhaps the most egregious example of this is when outlets report on patent filings, such as when the BBC reports that ‘A coffee-delivery drone that can detect when people are tired and bring them a drink has been patented by technology company IBM’ (23 Aug. 2018). Filing a patent requires neither that a company has a working prototype, nor even an intention of making one.

Importantly, not all industry-driven articles are positive. That AI-containing products are ‘creepy’ is a persistent concern – and outlets describe everything from Facebook’s eye-opening algorithms, to the Robotdog Spot Mini, to Chinese government surveillance as creepy. But perhaps more telling, a small fraction of articles also question the *ability* or *competence* of AI-containing products. Some argue that particular products simply don’t work well. The MailOnline and the *Daily Mail* published several feature stories meant to be sarcastic or humorous: one exclaims ‘Smart? These gadgets nearly drove me out of my mind’ (MailOnline, 7 Jan. 2018).

Other opinion pieces address more fundamental limitations of artificial intelligence itself. Some see that AI systems and humans will continue to have very different types of intelligence. Others argue that AI will continue to fail at tasks involving creative pursuits, emotional labour and relationships, and generating trust. Notably, however, these articles usually match this with the admission that AI already excels in many other areas. Others more directly question the fundamental idea of artificial intelligence itself. One commentator observes, ‘What is needed here is not artificial intelligence but real intelligence’ (HuffPost UK, 12 Feb. 2018). Others note that despite recent advances we remain very far from the long-running goal of a ‘general artificial intelligence’, a system that could replicate the full range of human intelligence (BBC, 17 Aug. 2018). That being said, such discussions of the fundamental limitations of AI are comparatively rare across the corpus.

Economics and Geopolitics

While every outlet addresses issues of economics and geopolitics, the two more right-leaning outlets, the print version of the *Daily Mail* and the *Telegraph*,

emphasise these concerns to a degree that sets them apart.

There is a degree of consensus that the UK is already a ‘world leader in AI’ (*Telegraph*, 20 Aug. 2018) that ‘punches well above its weight in AI’ (*Telegraph*, 26 Apr. 2018). Some echo the House of Lords in observing that ‘Britain contains leading AI companies, a dynamic academic research culture, a vigorous start-up ecosystem and a constellation of legal, ethical, financial and linguistic strengths located in proximity to each other’ (quoted in the *Telegraph*, 18 Apr. 2018). Others connect the UK’s success in AI to something more intrinsic and unique to the UK as a country, noting the UK’s ‘rich history in AI’ going back to Charles Babbage and Ada Lovelace (*Telegraph*, 26 Apr. 2018), or that, as the conservative MP Greg Clark argues in an editorial in the HuffPost, ‘We are a nation of innovators with some of the most brilliant minds and pioneering anywhere in the world’ (22 May 2018).

Two broad conversations predominate across discussions of economics and geopolitics in these outlets.

AUTOMATION AND THE FOURTH INDUSTRIAL REVOLUTION

As might be expected, discussion of automation and job loss is a persistent concern across coverage. There are several distinct ways outlets cover this topic. The left-leaning outlets frequently discuss jobs lost through automation. Some focus more on manufacturing jobs or those ‘in the lower-skill, lower-pay sections of our economy where job insecurity is already the new normal’ (HuffPost UK, 12 July 2018); others write more broadly about the impact across sectors. Several commentators in the *Guardian* consider the larger economic or social implications of massive job loss; one observes, ‘if robots are taking human jobs, we need to figure out how we would deal with a large jobless population’ (13 Mar. 2018).

In contrast, others argue that as AI takes over some jobs, new jobs will arise. Some go even further and claim that AI will be a net *producer* of jobs. This sentiment is found most commonly in the *Telegraph* and the *Daily Mail*. Jeremy Warner argues in the *Telegraph*, ‘We don’t need to worry too much about the long-term effects [of automation]. Anything that drives productivity growth, the magic ingredient that feeds prosperity and living standards, is by definition always going to be our friend’ (18 Apr. 2018).

In many ways, this argument is closely tied to the persistent discussion of AI as part of a larger ‘fourth

industrial revolution'. In its broadest form, the fourth Industrial Revolution is described as a revolution in economic and industrial life defined by 'advancing digital technologies, robotics and artificial intelligence (AI)' (BBC, 18 Apr. 2018) that will 'make many jobs obsolete with far-reaching social and cultural consequences' (*Telegraph*, 20 Aug. 2018). As indicated by its name, it is counterpoised against previous industrial revolutions each of which 'had a wrenching and lengthy impact on the jobs market, on the lives and livelihoods of large swathes of society' (Andy Haldane quoted in the *Telegraph*, 21 Aug. 2018). Yet, comparison to these previous transformations ultimately suggests an optimistic outlook: as one *Telegraph* editorial observes, 'each of the great technological advances of the past have enriched the world not impoverished it, so why should it be any different this time?' (22 Aug. 2018).

NATIONAL SECURITY AND REGULATION

Beyond motivating a transformation in the global economy, some articles frame AI as heralding a new geopolitical order. Five different articles, a mix of news, features, and opinion pieces, quote Putin as saying in 2017: 'Artificial intelligence is the future, not only for Russia, but for all humankind. Whoever becomes the leader in this sphere will become the ruler of the world.' While it remains unclear how this might happen, the implication is that whichever country leads in AI development will not only have economic advantage, but will have a *military* advantage both in terms of cyber- and traditional warfare. Several articles in the *Telegraph* detail arguments that the UK needs to increase military spending to update its military equipment and infrastructure and increase its capacity to fight cyber attacks.

For others, the imperatives of national security and prosperity mean we must limit regulation. Jeremy Warner argues in the *Telegraph* that 'We risk leaving the future to China in our rush to data protection' – that even while privacy concerns are important, 'data is the future' (25 Apr. 2018) and limiting the ability of companies to gather and use large data sets will set the UK back. Several other *Telegraph* opinion pieces similarly argue that the government 'must resist the temptation to slow down its [AI's] advance through inappropriate red tape and burdensome taxation' (23 Jan. 2018; see also 1 June 2018). Indeed, this imperative against regulation and taxation also means some are sanguine about AI's prospects post-Brexit. As a profile of Boris Johnson in the *Telegraph* notes: 'To him, the whole point of Brexit is to break free from EU regulations so Britain can lead the world in life sciences, artificial intelligence and driverless cars – all without hindrance from Brussels diktats' (8 June 2018). That is, once free

from restrictive EU regulations, the UK will be able to fully pursue AI research – and have unfettered access to lifeblood of AI: large amounts of *data*.

Ethics, Discrimination, and Killer Robots

Broadly speaking, the ethics of artificial intelligence is one of the most common themes across the corpus. Yet, just as the importance of AI to national success and security is prioritised in right-leaning papers, left-leaning outlets show a greater emphasis on the ethics, limits, and dangers of AI. Articles highlight ethical concerns surrounding topics such as deepfakes, automation, autonomous vehicles and weapons, data, privacy, hiring, facial recognition, human enhancement, and discrimination.

Just as industry products and initiatives drive much of the coverage of AI more broadly, they also serve as news pegs for many considerations of ethical issues. Ethical discussions appear in stories as diverse as those about a special effects company using AI or Google's 'ethically lost' Duplex (Zyner Tufekci quoted in the *Guardian*, 11 May 2018). Similarly, there are a series of news and feature articles about ethics initiatives at tech companies, and many articles that report high-level industry executives advocating for more ethical implementation of AI. This includes Demis Hassabis of DeepMind, Matt Wood at Amazon, and Kaave Pour at IKEA, who is cited saying 'We truly believe that it's going to be a competitive advantage to behave ethically' (Wired UK, 20 Mar. 2018). Of course, the most commonly cited CEO is Elon Musk, whose concerns about AI taking over the world appear in 88 different articles, nearly 12 percent of the entire corpus. Of these, more than half are in the MailOnline.

While there are a handful of opinion pieces across outlets that present sophisticated and rich discussions of the ethics of AI, many more articles substitute calls for discussion about ethics for actual ethical discussion. Articles frequently identify ethical *topics* or *questions*, but then stop before going further. One *Guardian* article notes, 'Whenever there is talk of enhancing humans, moral questions remain – particularly around where the human ends and the machine begins' (1 Jan. 2018). A HuffPost editorial observes,

The real issue is whether we really want a society where a piece of code decides which information sources we do and do not see? An even better question might be, do we really want to live in a society where that piece

of code is owned, built and operated by a private corporation? (13 July 2018)

Sometimes this questioning without answering involves pushing the work of actual ethics off on others, such as academics, government, or one of a number of new organisations meant to address the ethics of AI, including the UK government's new centre on Data and Ethics, Deepmind's Ethics & Society, and the Partnership on AI.

While, as noted above, there are a number of different ethical topics raised across the corpus, two of the most prominent involve discrimination and autonomous weapons – or 'killer robots.'

DISCRIMINATION

Both the *Guardian* and the *HuffPost* share a persistent concern that both current AI systems and those in development can enact and perpetuate forms of discrimination. Articles provide a range of examples of AI systems that have already shown significant biases. One *Guardian* article reports that after Google was widely criticised for an image recognition system that 'auto-tagged pictures of black people as "gorillas"' Google's solution was to 'prevent Google Photos from ever labeling any image as a gorilla, chimpanzee, or monkey – even pictures of the primates themselves' (12 Jan. 2018). Several news and feature articles report on bias in AI systems used for hiring. A *Wired* piece reports on AI-containing systems used by police to discriminate against lower-income neighbourhoods (1 Mar. 2018).

More telling however, are the *explanations* of why AI is discriminatory – of where the bias itself comes into the system. Articles identify at least three distinct source of bias. First, AI systems require large data sets. Four different articles quote 'that the old computing adage "garbage in, garbage out"' – that is, 'Bias creeps in when your data sets aren't inclusive enough and AI then learns from our own prejudices' (*HuffPost* UK, 12 Mar. 2018).

Second, others ask, as one BBC article does, 'what if the algorithms themselves are biased?' (6 Feb. 2018). This is especially troubling given 'a lack of transparency about what goes into the algorithms' (*Wired* UK, 24 Aug. 2018). On the one hand, machine learning processes can make it impossible to understand the ways in which algorithms make decisions. On the other, commercial interests can mean companies are not willing to reveal how algorithms work. One op-ed in *Wired* identifies a unique form of algorithmic bias in the 'inbuilt tendency to favour that which can be measured over that which cannot' (28 Jan. 2018).

Third, some note that bias in AI systems can derive from the engineers who construct them. One BBC article quotes a computer scientist observing, 'sexist AI could be down to the fact that a lot of machines are programmed by "white, single guys from California" and can be addressed, at least partially, by diversifying the workforce' (BBC, 6 Feb. 2018) or by 'encouraging more women to take up the profession and create algorithms' (*Guardian*, 13 Mar. 2018).

KILLER ROBOTS

Notwithstanding its cartoonish moniker, the discussion over 'killer robots' in the corpus is for the most part grounded in a pressing concern over the use of autonomous or semi-autonomous weapon systems. Despite the occasional sensationalist headlines in the *MailOnline*, such as 'Killer military robots will create a nightmare dystopia if they are allowed to kill at will...' (27 Aug. 2018), much of the coverage hews close to grounded events and actions within the broader social action around autonomous weapons. A number of articles highlight the work of the Campaign to Stop Killer Robots, the recent UN hearings on autonomous weapons, the Future of Life Pledge – a statement against autonomous weapons signed by more than 2,500 researchers and industry experts – and the organised opposition to Google's contract to develop automated drone detection software for the US Department of Defense.

Interestingly, there is little discussion in the corpus whether autonomous weapons should be banned. Rather, what discussion or disagreement exists concerns what are the specific dangers of autonomous weapons. Articles identify three distinct dangers of autonomous weapons.

The first danger, which is found more in right-leaning outlets, is an extreme or 'existential' threat such as that autonomous weapons rebel. One *Telegraph* article begins: 'Scenarios from *The Terminator* [sic] in which beings with artificial intelligence turn on humans are just "one to two decades away", according to a former Google chief' (2 Mar. 2018), another suggests the possibility that AI 'becomes self-aware and attempts to wipe out humanity' (16 Apr. 2018).

In contrast, Toby Walsh of the *Guardian* writes:

The killer robots I'm talking about aren't T101 Terminator robots. It's stupid AI that I'm most worried about. They are much simpler technologies that are just a few years away. (6 Apr. 2018)

Walsh joins other commentators in noting that weapons will be ‘Cheap. Effective. And easily available,’ (ibid.) making these ‘weapons of mass destruction’ that ‘will industrialise war, changing the speed and duration of how we can fight’ (*Guardian* 9 Apr. 2018). Others worry that autonomous weapons are currently against international law, or that the race to create these systems could spark a new global arms race, or an actual war. Finally, some raise more basic moral or ethical concerns about ceding decisions over life and death to algorithms. The Campaign to Stop Killer Robots writes on its website: ‘Allowing life or death decisions to be made by machines crosses a fundamental moral line. Autonomous robots would lack human judgment and the ability to understand context.’

Understanding Media Coverage of AI

While there have been few empirical studies of UK media coverage of artificial intelligence, as noted above, some have described coverage as consistently sensational (Select Committee on AI, 2018; Schwartz, 2018). While we found some sensationalised content, we saw far less than expected. A number of MailOnline articles include sensational headlines, such as ‘The AI that can tell you when you’ll DIE...’ (23 Feb. 2018) – however, the articles that follow are far more routine, deriving much of their substance from press releases or other news articles. Some headlines across other outlets veer toward the sensational – often slightly overstating findings: ‘DeepMind has trained an AI to unlock the mysteries of your brain’ (Wired UK, 9 May, 2018). But most are far more routine: ‘Fintech firm Previser targets late-payment problem’ (BBC, 8 Jan. 2018) or ‘Move over CPUs and GPUs, the Intelligence Processing Unit is the super-smart chip of the future’ (Wired UK, 25 June 2018).

Our findings, however, reveal an alternative set of concerns regarding the coverage of AI. Irrespective of topic, a majority of articles across the corpus are pegged to industry concerns, products, and initiatives. Of course, much of the research and development of AI is occurring in the commercial sector. However, we identify several concerns with this persistent indexing to industry initiatives. When articles profile a new AI start up or detail some high-level business deal, largely on the basis of industry sources, they amplify self-interested assertions of AI’s value and potential. Similarly, in positioning AI primarily as a private commercial concern, outlets undercut consideration of the role of politics, public action, and collective decision-making in addressing AI. In prioritising

industry sources above government employees, politicians, activists, and academics, outlets downplay the responsibility of publics and public representatives in addressing this emerging public issue. At the same time, in being so consistently indexed to industry sources and concerns, outlets limit the range of voices included in public conversation. For example, every time that an article reports Elon Musk’s extreme – if entertaining – opinions about AI, it misses an opportunity to bring in other, less familiar voices.

Our analysis also shows how industry topics and sources regularly encourage outlets to position AI as a viable solution to a wide range of problems. When articles describe AI systems that can direct coffee-dispensing drones, identify clothing brands from pictures, reorder geopolitical power, or conquer death, the implication is AI can serve as a solution to a vast array of problems – from the frivolous to the profound. Yet in doing so, outlets rarely interrogate either the limits of AI’s competency or the role that humans continue to play in its design and implementation. As relevant, competent, and somewhat autonomous, AI is often described as a radical disrupter, up-ending the economic and political status quo. As such, AI is seen as already wielding massive influence across our lives: reshaping everything from global economics, to politics, to healthcare. Even (calls for) ethical considerations of AI assert the revolutionary effects of AI. It is only *because* AI will have such profound and radical influence across sectors that we need to consider its ethical implications. The website for DeepMind’s Ethics & Society initiative announces ‘We created DeepMind Ethics & Society because we believe AI can be of extraordinary benefit to the world, but only if held to the highest ethical standards’ (DeepMind, 2018). While it is possible that AI will radically re-order all areas of our lives, AI experts continue to disagree sharply about AI’s impact (e.g. House of Lords, 2017) and the degree to which social action, politics, and public decision-making will amplify, impede, or mediate AI’s effects.

Our findings also highlight the connections between AI coverage and technology reporting practices amid structural changes seen across health, science, and technology journalism over the past two decades. Existing research has already shown how financial pressures have both encouraged organisations to cut back specialty desks and undercut financial resources needed to complete in-depth and/or investigative articles (Schäfer, 2017). In response, many outlets have come to rely heavily on press releases for day-to-day science and technology news stories (Brennen, 2018; Lynch et al., 2014), a trend seen across our corpus as

well. While academic and government organisations now regularly produce and distribute press releases, industry has something of a leg up. For many industry events or releases, there are often not only press releases, but entire media campaigns. For example, when IBM held its 'debate' between an AI system and two professional (human) debaters, the company not only brought together journalists and project researchers, it also flew out experts not involved in the project to serve as 'independent' sources for news stories (Anon, personal communication, 24 Oct. 2018). The director of research at IBM, Arvind Krishna, also published a blog post about the event in which he opined broadly on the contribution IBM had made. While the event itself was meant to generate media coverage, in providing journalists with story formats, narratives, sources, and quotes, IBM offered journalists all the materials necessary for easy-to-write stories.

Finally, our analysis also reveals the beginnings of a politicisation or polarisation in the way that AI is covered. Rather than sides of a single issue, this politicisation concerns the topics that different outlets see as constituting artificial intelligence itself. Even so, it is important not to overstate this incipient politicisation. All six outlets cover a range of topics, and, as shown above, there remain notable similarities in the ways they cover AI. The key difference lies in the topics outlets choose to emphasise. We have shown how right-leaning outlets consider AI through questions of economics, business, and national security – topics long prioritised by the Conservative party. Similarly, left-leaning outlets emphasise questions of ethics, discrimination, and privacy – reading AI through long standing concerns over labour relations and social justice. This growing politicisation is at least partly a function of the UK's politically divided news landscape (Newman et al., 2018) in which outlets have long held strong political affiliations or sympathies. When faced with a new and complicated topic, some outlets seem to be considering AI through existing political issues and frames. While these approaches might help outlets order and make sense of this emerging technology, they can also exacerbate a more general fracturing of public conversation. In a much larger sense, these findings reveal that rather than a single public conversation, there are many distinct discussions occurring around artificial intelligence. Beyond political fracturing, some people think AI will help us. Others think it will kill us. While a diverse landscape of conversation is, arguably, productive, these different discussions are rarely connected in news coverage. Ultimately, is not only that different outlets emphasise different topics when discussing AI; it is

that those outlets *define* in different ways what AI is as a public issue.

This fracturing in the topics that define AI is further demonstrated by a handful of articles in the *Guardian* that argue that there are dangers in outlets focusing on the wrong topics. These pieces observe that media have focused far more on the sensational but unlikely existential threats of AI and so have failed to address the far more real and pressing dangers or issues. Beyond the charge of sensationalism, this claim relies on an assumption that there is a limited attention economy such that as articles discuss the possibility of AI taking over the world, they are unable to participate in a more grounded and useful conversation. This is, ultimately, a disagreement over what topics should give shape to AI as a public issue.

Conclusions

News coverage provides an important basis for public discussion of AI, and as the issues surrounding these technologies grow more important, it is worth considering how journalistic treatment might evolve.

First, our finding that industry sources dominate coverage suggests the importance of including a wider range of voices in discussions of AI. Academics, activists, politicians, civilians, and civil servants, amongst others, can all contribute to a rich and sophisticated public debate around AI. It is not only important that journalists seek out diverse voices, but also that they actively join the discussion of these issues of common concern to complement perspectives coming out of industry and the private sector.

Second, precisely because AI is likely to have extensive and profound implications across our societies, it is important for news outlets to explicitly recognise the legitimately different political interpretations of what AI is and ought to be. As part of an effort to cover the society-wide implications and diverse possibilities of AI, outlets could also prioritise more collaboration across news desks. As AI develops as a public issue, it is necessary to interrogate its relation to many other realms beyond technology, including politics, economics, and health. Collaborations between journalists on different beats could help outlets produce more sophisticated articles about AI, allowing journalists to bring together their varied expertise and source contacts.

Ultimately, these recommendations suggest that outlets should avoid uncritically considering AI on the

basis of industry sources. Outlets have largely done a good job at showing both the potential benefits and dangers of AI. However, some have struggled to show the strengths and the weaknesses of AI as applied across sectors or to position AI as a fully public problem, one that requires a diversity of voices to address. While AI can do some things impressively well, it is not a solution to every problem. Similarly, despite what Elon Musk, Stephen Hawking, or Vladimir Putin say, the wider public implications of AI remain unclear. While media should explore in detail the promise and pitfalls of AI, they would be well served to treat it less as a world-shaking revolution and more as a set of technologies in the process of being designed, a set of choices in the process of being made, and a set of problems in the process of being collectively solved.

References

- Brennen, J. S. 2018. *Science In Pieces: Public Science in the Deformation Age* (Doctoral dissertation). Retrieved from Proquest Dissertations Publishing, (Ref 10787597).
- Craig, C. 2018. 'Conflicting Truths. How Does Government Listen to Scientists?' Oxford Martin School.
- Dafoe, A. 2018. *AI Governance: A Research Agenda*. Oxford: Future of Humanity Institute.
- Dunwoody, S. 2014. 'Science Journalism', in *Routledge Handbook of Public Communication of Science and Technology* 2nd Edition, Bucchi M., and Trench B., (eds), 27-39. London; New York: Routledge.
- Gov.uk. 2018. 'The Grand Challenges.' <https://www.gov.uk/government/publications/industrial-strategy-the-grand-challenges/industrial-strategy-the-grand-challenges#artificial-intelligence-and-data>
- Kueng, L. 2017. *Going Digital. A Roadmap for Organisational Transformation*. Oxford: Reuters Institute for the Study of Journalism.
- Lynch, J., Bennett, D., Luntz, A., Toy, C., and VanBenschoten, E. (2014). 'Bridging Science and Journalism Identifying the Role of Public Relations in the Construction and Circulation of Stem Cell Research Among Lay People', *Science Communication*, 36(4): 479-501.
- Newman, N., Fletcher, R., Kalogeropoulos, A., Levy, D. A. L., and Nielsen, R. K. (2018). *Reuters Institute Digital News Report 2018*. Oxford: Reuters Institute for the Study of Journalism.
- Schäfer, M. S. 2017. 'How Changing Media Structures Are Affecting Science News Coverage', in *The Oxford Handbook of the Science of Science Communication*, Hall Jamieson, K., Kahan, D., and Scheufele, D. A. (eds), 51-60. New York: Oxford University Press.
- Select Committee on Artificial Intelligence. 2017. Corrected Oral Evidence. House of Lords.
- Select Committee on Artificial Intelligence. 2018. *AI in the UK: Ready, Willing and Able?* UK: House of Lords.

Acknowledgements

The authors would like to thank Lucas Graves and Liliana Bounegru for their help and guidance during the course of this research, as well as the rest of research team at the Reuters Institute for the Study of Journalism and the Computation Propaganda Project at the Oxford Internet Institute.

ABOUT THE AUTHORS

J. Scott Brennen is a Research Fellow at the Reuters Institute for the Study of Journalism and the Oxford Internet Institute at the University of Oxford.

Philip N. Howard is the Director of the Oxford Internet Institute and a Professor of Sociology, Information and International Affairs.

Rasmus Kleis Nielsen is the Director of the Reuters Institute for the Study of Journalism and Professor of Political Communication at the University of Oxford.

Published by the Reuters Institute for the Study of Journalism as part of the Oxford Martin Programme on Misinformation, Science and Media, a three-year research collaboration between the Reuters Institute, the Oxford Internet Institute, and the Oxford Martin School.