



Abstract

In debates about artificial intelligence (AI), imaginations often run wild. Policy-makers, opinion leaders, and the public tend to believe that AI is already an immensely powerful universal technology, limitless in its possibilities. However, while machine learning (ML), the principal computer science tool underlying today's AI breakthroughs, is indeed powerful, ML is fundamentally a form of context-dependent statistical inference and as such has its limits. Specifically, because ML relies on correlations between inputs and outputs or emergent clustering in training data, today's AI systems can only be applied in well-specified problem domains, still lacking the context-sensitivity of a typical toddler or house-pet. Consequently, instead of constructing policies to govern artificial general intelligence (AGI), decision-makers should focus on the distinctive and powerful problems posed by narrow AI, including misconceived benefits and the distribution of benefits, autonomous weapons, and bias in algorithms. AI governance, at least for now, is about managing those who create and deploy AI systems, and supporting the safe and beneficial application of AI to narrow, well-defined problem domains.

Specific implications of our discussion are as follows:

-



Acknowledgements

This research was funded in part by the Ewing Marion Kauffman Foundation with additional support from Center for Information Technology Research in the Interest of Society (CITRIS), The Trigger Project of the European Commission, and the Research Institute for the Finnish Economy (ETLA). Particular thanks to Tim Büthe, professor and chair for International Relations at the Hochschule für Politik/School of Governance and School of Management at the Technical University of Munich. This paper was sparked by John Zysman's long-standing collaboration with him on technology governance topics, as well as our participation in the TUM Workshop on Governance of New Technologies, organized by Tim Büthe and Hanna Hottenrott in February 2020, and in a special issue on AI governance of the Journal of European Public Policy, co-edited by Tim Büthe with Christian Djefal, Sabine Maasen, and Nora von Ingersleben-Seip, to appear in 2021. The contents of this publication are solely the responsibility of the authors.

Note: A draft of this paper was previously posted as a BRIE Working Paper #2020-5 in August 2020



Introduction	2
Part 1: A framework for discussing AI: What it is and what it is not	4
What AI is not (or at least not yet)	4
AI as a system of statistical inference	5
What can today's AI systems do?	6
What are the risks associated with these capabilities?	6
What are the limitations of today's AI systems?	7
What is going on here? Interpreting reality	9
Part 2: AI in the Community: Framing the Governance Problem	10
Who governs AI?	11
What should be governed?	11
Part 3: AI, Competition, Growth, and Jobs	14
Investing in AI for growth and power	14
Will AI accelerate growth?	16
AI, distribution, and the impact on the workforce	17
Part 4. Turning the Governance Story Around: AI in Government and Policy	18
Conclusion: Implementing Governance	19



Understanding the Limits, Possibilities, and Risks of AI in an Era of Intelligent Tools and Systems

John Zysman¹ and Mark Nitzberg²

August 2020

“...we need to reconceptualize it [AI/Deep learning]: not as a universal solvent, but simply as one tool among many, a power screwdriver in a world in which we also need hammers, wrenches, and pliers, not to mention chisels and drills, voltmeters, logic probes, and oscilloscopes. In perceptual classification, where vast amounts of data are available, deep learning is a valuable tool; in other, richer cognitive domains, it is often far less satisfactory.”

- Gary Markus (p. 18)³

“Intelligent tools are diffusing through our economies and society. Some of the developments are powerfully changing how our economies work and how we live our lives. Some of the purported developments are simply hype. Amid the froth, many believe that our current social and economic arrangements will be swept aside Others, ourselves included, assert that the world is ours to create. That is easy to assert but difficult to demonstrate and harder still to implement.”

- John Zysman, Martin Kenney, Laura Tyson (p. 2)



Introduction

Governance in an era of intelligent tools and systems requires stepping beyond the hype about particular tools and the despair that claims regarding their potential can engender. The phrase “intelligent tools and systems” points, therefore, to the toolbox, not individual tools. The toolbox itself is constantly expanding, and the tools are constantly gaining power. Importantly, given the suite of intelligent tools, governance issues about particular tools cannot usefully be dealt with entirely in isolation. Hence, the debate about artificial intelligence (AI) governance must be about the set of intelligent tools and how they relate to each other.

Consider what is left out of the conversation if we only discuss AI. Take, for example, two-sided digital platforms (hereafter termed digital platforms). Digital platforms, key tools in the box, are usually treated separately from AI and big data. But digital platforms are very much the nexus of both. Platforms amass the big chunks of data required to effectively deploy AI tools. In turn, AI tools, which depend on big data, give power to the creators of these platforms.

Consequently, digital platforms, one may argue, are the most fundamental of the tools in the tool box, often creating the framework for developing big data and applying statistical tools such as deep learning to that data.

Platforms “are an emblem and embodiment of the digital era just as factories were of the industrial revolution” (Bearson, et al., 2019; Kenney, et al., 2019). Thus, focusing only on AI would be like focusing only on steam engines and the rules for steam engines when factories first emerged. Of course, digital platforms, data, and AI applications all run on and depend on the ever-increasing power of digital infrastructure and computing in the cloud. So, governing platforms, a significant debate in itself, is entangled within any discussion about governing data and AI. Seen from that vantage, AI is simply a part of the challenge of governing the platform economy—the economy of intelligent tools and systems.

To clarify our position, the core story about governance in a digital era is not about the rules of particular tools, but the challenges represented by a new tool box. That toolbox includes big data, powerful computing capacity, algorithms, and software in general. AI, in its several current manifestations, is one tool amongst many, a truly powerful tool with an ever-expanding set of applications.⁴ The most valuable firms in the world are increasingly built using the whole toolbox, not just one tool (McGee & Chazan, 2020; Kenney, 2020).



New technologies always present possibilities and risks: a reminder from the railroad era that balancing possibilities and risks is not simple. The 1895 Gare Montparnasse train derailment exemplifies how we can get caught up in the boundless possibilities of a new technology, but fail to fully work out safety and governance concerns.



The focus of this paper, nonetheless, is governance of the array of digital tools loosely labeled AI. Discussion about AI currently focuses on deep learning and machine learning (ML), but indeed deep learning and AI are simply the latest tools, labeled 'AI', in the toolbox.

There are, in our view, two separate debates about AI:

- Community and social preferences: One debate, about matters such as privacy and discrimination in applications, is really about molding AI usage to our community and social preferences. The array of discussions about AI ethics certainly falls here.
- Economic and strategic advantage: A second debate is whether, and for whom, AI creates economic and strategic advantage and how best to promote the development and competitive deployment of AI for economic, geo-strategic, and military advantage.

These two discussions often collide.



Part 1: A framework for discussing AI: What it is and what it is not

The fourscore-year history of AI and its many definitions preface so many articles, trade books, and documentaries these days that we indulge only a few succinct comments on the history. As programmable computing machines emerged in the 1940's, British logician Alan Turing contemplated "intelligent machinery that could learn from experience ... by altering its own instructions." In 1956, the discipline was born and named AI at the "Dartmouth Summer Research Project on Artificial Intelligence" (Leslie, 2019; Russell, 2019).⁵ It has since experienced surges of public interest and funding (ca. 1956-1974, ca. 1981-1992), each followed by a so-called "AI winter" when funding and public interest were scarce. Around 2015, a new AI boom took hold and "ai" began replacing ".com" on buildings and billboards.

The field has been defined by two parallel aspirations:

- A foundational aspiration, "general AI" or "AGI," to build human-level intelligence in a computer program.
- Smaller, more or less well-defined aspirations, often characterized as AI's moving target ("AI is whatever we haven't solved yet"). These goals have included playing chess, transcribing speech, recognizing objects in a picture, diagnosing illnesses from a set of symptoms, answering questions based on a given text, proving theorems, and detecting nuclear tests in massive seismology data. These narrower aspirations are called "narrow AI."

Technologies developed by AI researchers had become commonplace in many systems but were not called 'AI' even in the mid-2000's. Examples include speech recognition, image processing, data mining, industrial robotics, medical diagnosis, search engines, and recommendation systems for news, books, and films. Their performance varied and, although highly valuable commercially, most performed below human-level.

As in all articles, we round the bend to discuss the power of deep learning. Around 2010, deep neural networks and related machine learning techniques began showing extraordinary results, many that approached or exceeded human-level performance in areas such as speech transcription, face recognition, and medical image diagnostics. The anthropomorphic aspect of AGI seemed to come within reach; ignited a global AI "arms race;" drove Microsoft, Google, IBM, and many others to become "AI-first" companies; and seeded a wide range of efforts intended to assure that this newly powerful AI is safe and beneficial, and, certainly, to reassure customers and the public.

What AI is (and is not)

Personal digital assistants Alexa and Siri (named for marketing effect) respond sensibly to many spoken requests; OpenAI's GPT3 "writes" pages of coherent text in a consistent style based solely on a two-line prompt; Google Lens identifies plants and architectural landmarks from a single photo; and Waymo's autonomous cars safely navigate difficult construction zones on the streets of San Francisco. How far off is technology that can think, in general, as well as a human?



There are two problems with this leap. The first lies in the words, “in general.” While each astonishing achievement of the last 10 years shows the power of the new AI clockwork, each is still a tool—and thus, narrow AI. This is the essence of AI’s importance these last 10 years.

The second problem is that, in colloquial use, the term AI implies the powerful aspiration to create a simulated human with the full spectrum of self-awareness, context-appropriate behavior, emotion, empathy, and the common sense about the world exhibited by a typical toddler or house pet.⁶



as the game score, that enables a system to learn by experience. Then, in operation, the system calculates the most likely or best output for a given set of inputs based on this "training."

In today's ML focused AI, the very nature of statistical inference at the core of deep networks confers limits. As mere reflections of the correlations between inputs and outputs in data, today's most well-known AI systems are inherently conservative; they operate poorly outside the space of the data they have already seen. An AI system can transcribe a person's speech because it has seen the correct strings of words paired with the same or very similar vocal sounds; a different AI system can identify a dog in a photograph whose arrangement of



One example is the AI systems employed by social networks, whose effectiveness in keeping users interested can end up changing what they are interested in. By providing consistent and attentive viewership from users, YouTube is highly appealing to advertisers. As such, YouTube's AI systems are tuned to the narrow objective of keeping users interested in each successive video in their automated queue. But as Professor Zeynep Tufekci of the University of North Carolina has observed, one of the algorithms' primary strategies for keeping users engaged is to show them ever more extreme versions of similar content (Friedersdorf, 2018). This can be as innocuous as showing videos about running ultramarathons after videos about jogging; or as potentially dangerous as videos of Donald Trump rallies lead to white supremacist rants and Holocaust denials. Performed at YouTube's immense scale, this strategy of engagement could alter how people think. A powerful statistical algorithm with a narrow objective has, thus, broad social consequences that by definition could not be incorporated into its reward model.

Another example of AI's unintended consequences beyond the fulfillment of narrow objectives lies in the facial recognition company Clearview AI. With a library of over three billion images, Clearview AI empowers customers to upload a photo of an individual and in response learn the identity and any discoverable information on the internet about that person (Hill, 2020). A New York Times investigation in 2019 found that Clearview AI had been adopted by over 600 law enforcement agencies worldwide, with little to no scrutiny. With its massive library of data, its algorithm is remarkably effective at the helping security services identify individuals. But the very scale underlying the system's effectiveness has resulted in a significant impact on the much broader domain of personal privacy.

Broader unintended consequences are inherent to the combination of narrow objectives and significant scale. Efficiency is a limited metric when evaluating the pursuit of goals by powerful entities; but efficiency is the narrow objective that most of today's systems are optimized to achieve. This makes the risks of today's AI difficult to mitigate, and suggests that greater caution is needed in deploying systems at scale.

W |  AI ?⁶



“at times responsible for the composition, coloring and other aspects of a work” (Cohen, 2016). However, its creations depended heavily on randomization (Deutsch, 2016).

Some of AI's limits can be broadly indicated by the inability to generalize. Human children, as has been argued by Professor Allison Gopnik of UC Berkeley, develop general frames or models of the world, shuffling amongst models to learn and interpret what is going on around them (Gopnik, 2011; Samuel, 2020). AI systems have no such frames or models to refer to. As Gary Marcus has written, AI systems do not know what they are for; nothing in their necessarily narrow data tells them anything about the purpose of their task, which depends on the world in which it sits (Marcus & Davis, 2019).

The inability to generalize prevents current AI systems from three forms of reasoning, which power humans' interaction with the world. First, AI systems cannot yet reason about causality. By observing daily life around them, humans at a very young age understand how a physical action causes a result. But as has been documented by MIT Professor Joshua Tenenbaum, AI systems cannot answer basic causal questions about a scene, such as “what caused the ball to collide with the cube?” (Knight, 2020). A machine may detect a pattern in image pixels that represents a ball colliding with a cube, but has no general model of the world telling it that the ball moved because a person pushed it. This has obvious implications for AI systems intended to function in a physical environment, as with the autonomous vehicle systems noted above. But a baby's ability to reason causally about the physical world underlies their ability to reason about cause and effect in the abstract as they get older. This, in turn, is a core factor in human's ability to be effective in situations that have not been encountered before (Knight, 2020). So long as AI's reasoning is confined to detecting correlations among arrangements of pixels, words, or other superficial data, it will not be able to work through unfamiliar situations.

A machine may detect a pattern in image pixels that represents a ball colliding with a cube, but has no general model of the world telling it that the ball moved because a person pushed it. This has obvious implications for AI systems intended to function in a physical environment

Second, AI systems cannot yet understand human emotion. Human emotional understanding is rooted in an ability to know that one feels what another feels, and in a shared grounding in the human condition. Neurons indicating the ability to signal “I feel what you feel,” to empathize, have been observed in social animals beyond humans, including chimpanzees (Blakeslee, 2006). But such interactions can only be simulated in AI systems, which have no background understanding or sense of shared experience. Additionally, simulated understanding of emotion is in its infancy in AI research.

Third, AI systems are unable to use judgment. Judgment is the ability to step back from the immediate domain and understand something's significance to the broader world. Humans exercise judgment by relating an immediate event to their general model of the world and reason about how they are related to one another. But while today's AI systems may be able to detect something unusual in a pattern, they lack the general world model necessary to determine whether the appropriate response is “that's odd,” or “eureka!”



What is going on here? | Mr. Robot

Humans' distinctive aptitude for these three types of reason (identifying causation, understanding emotion, and exercising judgement and generalizing), can perhaps be assessed in terms of three core aspects of our cognition and interpretation: *context*, *narrative*, and *worldview*.

Context refers to how we answer the question, "what is going on here?"⁷ Mr. Robot certainly does not know.

In some sense, *context*—derived from our "models" of how the world works and our frames of reference to the world—defines for us the answers to the question, "what is going on here." In characterizing context, we are specifying which elements in a situation are relevant and which elements fall to the background. Of course, the aspects relevant to a context are continuously changing. For example, in high school, you and friends are being boisterous and disruptive before class. Then an adult walks in: is that person a teacher who demands attention, or a custodian who can be noticed and ignored? The answer to this question rests in how we define context, which constitutes our frames of reference for how the world works.

Seen thus, specification of context is a uniquely human capacity, linked to our experiences and the way we understand the world. That conclusion forces a question: in which applications can we apply AI tools without a human definition of context? In other words, when is a narrower specification of context sufficient? In which applications is a human definition of context essential?

Contexts do not stand alone. They are part of a sequence that constitutes the *narratives* of our lives. Defining context involves defining what "story" one is in at any given moment—what is the narrative? We live in socially defined narratives with the people around us, the rewards and threats we face, and much more. Different narratives imply different definitions of context; if how one defines context depends on his or her "narrative" about life, then, in similar situations, different people will define contexts in different ways.

Our narratives are themselves framed within, what are labeled in social science, "worldviews"⁸ Narratives about ourselves, others, and the world can be built into broader narratives about our communities and others, that constitute a view of how the world works as a whole. Religion, economic philosophy, or any number of broader narratives, worldviews, create the environment in which our individual narratives emerge.

Taken together, context, narrative, and worldview comprise much of what is distinctive about human intelligence, and as yet beyond the reach of statistics-based AI systems. Our continuous definition and redefinition of context, and ability to refer back and forth between it and the immediate objective, is essential to our ability to understand causality, emotion, and judgment and make essential generalizations.

As generalizability is a core aspect of AGI, this casts doubt on the possibility that systems built on today's dominant ML models can achieve AGI. To be sure, significant breakthroughs are needed.



Who should govern AI? With the creation of an Organization for Economic Co-operation and Development AI Policy Observatory and the proposal for a United Nations Panel on Artificial Intelligence, there is a swirling debate about how to govern AI.¹¹ This is not simply a matter of abstracted ethics, but of which public institutions and whose institutions are best situated for regulation and governance. Should the regulation be at the city level, as has often occurred with Uber? At the state or regional level, as with California legislation on privacy, which was inspired by EU privacy law? At the international level, as bargains amongst states?

Indeed, what should be undertaken by public authorities, and what by private self-regulation? Individual instances of calamities emerging from private self-regulation, such as the Boeing 737 Max crashes, underpin the importance of balancing public authorities and private actors in regulation.

Google's AI Principles are another example of private self-regulation; are the principles merely suggestions or does Google have a system for accountability? (Google, 2018) More generally, do public authorities require the detailed technical know-how embedded in private sector operations? If so, what should be the balance between private actors and public authorities as they collaborate on product regulation? Digital platforms represent the systemic problem of AI governance. Together, the algorithmic structure of the platforms and the "terms and conditions" to which those engaged on the platform agree make the digital platforms private regulators that are difficult to control by public authorities (Cutolo & Kenney, 2020; Kenney & Zysman, 2016).

Google's AI Principles are another example of private self-regulation; are the principles merely suggestions or does Google have a system for accountability? (Google, 2018) More generally, do public authorities require the detailed technical know-how embedded in private sector operations? If so, what should be the balance between private actors and public authorities as they collaborate on product regulation? Digital platforms represent the systemic problem of AI governance. Together, the algorithmic structure of the platforms and the "terms and conditions" to which those engaged on the platform agree make the digital platforms private regulators that are difficult to control by public authorities (Cutolo & Kenney, 2020; Kenney & Zysman, 2016).



that may be problematic in one domain may not have the same significance in another. A mistaken advertisement for shirts that draws the reader into a cycle of shirt advertisements is of less significance than targeted political misinformation or mistakes about medical diagnoses.

Not all crucial questions, however, can be identified or addressed in specific, narrow, contexts. Issues that reappear across sectors may be hidden or underestimated because the issue is not central in a particular domain. Some issues, such as bias, may be most evident or crucial in some domains, and most easily identified there, but are important throughout society.¹² Many issues that run across applications might have common solutions



We must also consider the actors themselves, the intent and purposes to which they put AI tools. That ultimately means governing the actors. Most evidently, we might consider regulating digital platforms as a means of regulating AI uses. Two-sided, digital platforms such as Google, Amazon, Facebook, Uber, and Yelp, amass data and apply these AI tools in diverse activities. Addressing the issues of AI applications in platforms will inevitably force us to consider issues of platform regulation.

Values and norms are really the issue. The obvious question, forced by an effort to govern AI, is what are the norms we want to support and encourage and the behaviors to be discouraged or banned? Indeed, the answers, outcomes of political conflict and social debate, do not simply pre-exist, awaiting implementation. Rather, the values, objectives, and preferred outcomes will be created by the process of creating a governance system.

If we focus on outcomes, the obvious issue is whether to consider general rules or domain specific concerns. As noted before, the types of data or tools that may be acceptable to make retail offers, or the scrutiny about those tools, may differ radically from the types of offers and scrutiny appropriate with financial offers or medical recommendations. Indeed, that is already the case: the Securities and Exchange Commission may review the financial offerings of Macy's but not their sales. So many of the issues must be debated and rules implemented, domain by domain. Moreover, some actors, such as digital platforms, cut across domains, and may require actor specific regulation of their deployment of AI tools. In some sense, it is a matter of how we want to derive general rules and figure out how to apply them across domains or how we want to infer general rules from sectoral domain situations. Certainly, there will be general principles, such as avoiding biased outcomes or limiting the extent of surveillance, that should frame particular domains or specific actor regulations.

In sum, AI raises new challenges and issues. Rules about AI cannot simply be grafted on pre-existing norms and regulations.

Part 3. AI, Competition, Growth, and Jobs

The implications of AI for economic growth and employment as well as for geopolitical competition is, for many, the central concern. Indeed, these concerns are, arguably, the drivers of public investment in AI.



Ought countries be investing massively to promote leadership in AI technology? Indeed, what would that





W/ AI ●● ↗ , ?




selling tractors to farm management services, different leadership skills will likely be required.¹⁸ The effective implementation of AI tools will likely move across a generation. The impact on growth is even less certain.

AI,

One element of the AI and growth story, the impact of AI-enabled automation on the workforce, is of particular concern to policy makers. Assume that firms effectively deploy intelligent tools, with AI firmly rooted in them, to win in the marketplace. What will happen to the workforce? The simple conclusion in this paper about governance is that the consequences of deployment for the workforce hinges on policy choices, not the technology per se.¹⁹

Sorting through the diverse claims about the impact of AI on jobs, on employment, is as difficult and problematic as sorting through the claims about growth. There have been competing claims that run from imminent disaster to slowly evolving disaster to perfectly manageable traditional reorganization of the economy. The extensive and diverse body of research²⁰ focuses variously on tasks that may be displaced, on jobs lost, and on sectors that may be changed. For the most part, the research does not focus on how these tools affect the organization, or better, the reorganization of work and responsibility, of how the activities, generally framed, of production, client or sales, or management can be reimagined and reconceived. Simplistically, we all know that jobs will be created, destroyed, modified, and transformed. What we don't know is how fast or what that transformation of work will look like. Indeed, there is a dispute about whether the deployment of these tools will down-skill the core of the economy, a dominant view at the moment, or create new possibilities and opportunities (Turca, 2016). There is a consensus view among economists on many aspects of the labor market issues: there is no evidence of long-term technological unemployment, productivity and employment growth go together, and the dislocation effects associated with changes in the demand for labor fall unevenly across workers, communities, occupations, and sectors (Tyson, 2019). In our view, the consequences will likely depend on the deployment strategies adopted by firms and how firms understand the market advantages these tools can create.²¹ In that sense, the labor market issues are wide open.²²

Part of an AI governance strategy should be to address the employment question. Encouraging the trajectory of the deployment of intelligent tools, all the more powerful with AI applications, to support the upgrading of existing work and the creation of new work is essential, and possible.²³

—  , government promotion of AI, a crucial part of governance, has at least two dimensions. One dimension is to encourage the development, diffusion, and use of the technologies. Setting aside purely military applications, achieving that goal requires assuring societal capacity to absorb and diffuse cutting-edge technology. The underlying science, military technology aside, will likely be widely available. Diffusion and application is essential, and that is a very different problem than investing in basic research or even development-oriented research. The diffusion question is rarely fully addressed.

The second dimension is the impact of AI on growth and employment. Strategies to capture the *growth possibilities* of intelligent tools in general, and AI in particular, require complementary development of a suite of complementary technologies and investment in a skilled workforce. Part of that is to promote business and deployment strategies that encourage complementarity and the upgrading of work and work opportunities.







... , AI applications are part of a suite of intelligent tools and systems that ultimately must be regulated as a set. Digital platforms, for example, generate the pools of big data on which AI tools operate. Thus, the regulation of digital platforms and of big data is part of the challenge of governing AI. Many of the platform offerings are, in fact, deployments of AI tools. Hence, focusing on AI alone distorts the governance problem.

... , the issues and choices will differ by sector. The consequences, for example, of bias and error will differ from a medical or a criminal justice domain to one of retail sales. Creating alternative mega-platforms for search or shopping may be quite difficult and breaking them up is often a pointless chimera. By contrast, alternative digital platforms for finance are being discussed as part of consideration of central bank digital currency (Bank for International Settlements, 2020).

... , the economic implications of AI applications are easily exaggerated. The possibilities for accelerating growth are overstated by would-be consultants, while the consequences for work depend less on the technology itself than how it is deployed in the reorganization of work in this era of intelligent tools and systems. Should public investment be concentrated on advancing basic research or on the diffusion of tools and the user interfaces and training needed to implement them?

...



Indeed, more likely than agreement on global governance, there will be intense cyber rivalries that risk splintering



16 One might more bluntly ask how the losers, those disrupted or displaced, may be compensated, ignored, or suppressed. For an argument about this in earlier economic restructurings see Zysman, 1983.

17



- » In turn, the organizational differences often turn on attitudes to the value of the workforce, what workers know that is crucial and indispensable.

Work we underscore, is being transformed in the digital era. But the debate is wide open about how fast, when, how much will be displacement and how extensive will be reorganization, or why we already observe distinct national variation in both use of the technology and its impact on job distribution and equality. That ambiguity suggests there is room for choice. The Northern European countries, which depend heavily on skilled labor and those countries, such as Japan, that are facing skilled labor shortages are hypothesized to provide the richest search domain for positive examples.

24 “There is no U.S.-based wireless access equipment provider today that builds those solutions,” said Sandra Rivera, a senior vice president at Intel who helps guide the chipmaker’s 5G strategy (Fung, 2019).

WORKS CITED

Andrews, D., Criscuolo, C., & Gal, P. (2016). The best versus the rest: The global productivity slowdown, divergence across firms and the role of public policy. *OECD Productivity Working Papers*, 5. <https://doi.org/10.1787/63629cc9-en>

Arntz, M., Gregory, T., & Zierahn, U. (2016). The risk of automation for jobs in OECD countries: A comparative analysis. *OECD Social, Employment and Migration Working Papers*, 189. <https://doi.org/https://doi.org/10.1787/5jlz9h56dvq7-en>

Bank for International Settlements. (2020). III. Central banks and payments in the digital era. *Annual Economic Report 2020*, 67-95.

Bearson, D., Kenney, M., & Zysman, J. (2019). Measuring the impacts of labor in the platform economy: New work created, old work reorganized and value reconfigured. *BRIE Working Paper Series*. <http://dx.doi.org/10.2139/ssrn.3363003>

Blakeslee, S. (2006, January 10). Cells that read minds. *The New York Times*. <https://www.nytimes.com/2006/01/10/science/cells-that-read-minds.html>

Brundage, M., Avin, S., Clark, J., Toner, H., Eckersley, P., Garfinkel, B., Dafoe, A., Scharre, P., Zeitoff, T., Filar, B., Anderson, H., Roff, H., Allen, G. C., Steinhardt, J., Flynn, C., hÉigartaigh, S. Ó., Beard, S., Belfield, H., Farquhar, S., Lyle, C., Crootof, R., Evans, O., Page, M., Bryson, J., Yampolskiy, R., & Amodei, D. (2018). The malicious use of artificial intelligence: Forecasting, prevention, and mitigation. arXiv. <https://arxiv.org/pdf/1802.07228.pdf>

Christensen, C. (1997). *The Innovator’s Dilemma: When New Technologies Cause Great Firms to Fail*. Boston: Harvard Business Review Press.

Cohen, Paul. (2016). Harold Cohen and AARON. *AI Magazine*, 37, 4, 63+.



the platform economy. *Academy of Management Perspectives*. <http://dx.doi.org/10.2139/ssrn.3372560>

Deutsch, M. (2016, July 8). Harry Potter: written by artificial intelligence. *Deep Writing*. <https://medium.com/deep-writing/harry-potter-written-by-artificial-intelligence-8a9431803da6#.p53bsfs3t>

Eckersley, P. & Nasser, Y. (2017). AI progress measurement. *Electronic Frontier Foundation AI Progress Measurement Project*.

Fjelland, R. (2020). Why general AI will not be realized. *Humanities and Social Science Communications*, 7(10), 1-9. <https://doi.org/10.1057/s41599-020-0494-4>

Friedersdorf, C. (2018, March 12). YouTube extremism and the long tail. *The Atlantic*. <https://www.theatlantic.com/politics/archive/2018/03/youtube-extremism-and-the-long-tail/555350/>



Koebler, J. (2015, March 23). Elon Musk on superintelligent robots: We'll be lucky if they enslave us as pets. *Vice*. <https://www.vice.com/en/article/nzekxq/elon-musk-on-superintelligent-robots-well-be-lucky-if-they-enslave-us-as-pets>

Lee, K. (2018). *AI Superpowers: China, Silicon Valley, and the New World Order*. Boston: Houghton Mifflin Harcourt.



- Samuel, S. (2020, February 28). Kids' brains may hold the secret to building better AI. *Vox*. <https://www.vox.com/future-perfect/2019/2/28/18239993/artificial-intelligence-children-machine-learning-alison-gopnik-psychology>
- Silva, S. & Kenney, M. (2019). Algorithms, platforms, and ethnic bias. *Communications of the Association of Computing Machinery*, 62(11), 37-39. <https://doi.org/10.1145/3318157>
- Smit, S., Tacke, T., Lund, S., Manyika, J., & Thiel, L. (2020). The future of work in Europe: Automation, workforce transitions, and the shifting geography of employment. *McKinsey Global Institute*. <https://www.mckinsey.com/featured-insights/future-of-work/the-future-of-work-in-europe>
- Turca, D. (2016, June 12). Is technology downskilling or upskilling us?. *Boostzone Institute*. <http://www.boostzone.fr/2016/06/12/is-technology-downskilling-or-upskilling-us/>
- Tyson, L. (2019). How automation will affect the future of work in Germany. *The Berlin Journal*. n.a.
- Zysman, J. (1983). *Governments Markets and Growth: Financial Systems and Politics of Industrial Change*. Ithaca, NY: Cornell University Press.
- Zysman, J., Feldman, S., Kushida, K., Murray, J., & Nielsen, N.C. (2013). Services with everything: The ICT-enabled digital transformation of services. In D. Breznitz and J. Zysman (Eds.), *The Third Globalization: Can Wealth Nations Stay Rich in the Twenty-First Century?*. Oxford: Oxford Scholarship Online, pp. 99-129.
- Zysman, J., Kenney, M., & Tyson, L. (2019). Beyond hype and despair: Developing healthy communities in an era of intelligent tools. *Innovation Policy White Paper Series 2019-01*. Munk School of Global Affairs, University of Toronto. <http://dx.doi.org/10.2139/ssrn.3414691>

