


Artificial intelligence, human intelligence and hybrid intelligence based on mutual augmentation

Big Data & Society
July–December: 1–6
© The Author(s) 2022
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/20539517221142824
journals.sagepub.com/home/bds


Mohammad Hossein Jarrahi¹ , Christoph Lutz² 
and Gemma Newlands²

Abstract

There is little consensus on what artificial intelligence (AI) systems may or may not embrace. Although this may point to multiplicity of interpretations and backgrounds, a lack of conceptual clarity could thwart the development of common ground around the concept among researchers, practitioners and users of AI and pave the way for misinterpretation and abuse of the concept. This article argues that one of the effective ways to delineate the concept of AI is to compare and contrast it with human intelligence. In doing so, the article broaches the unique capabilities of humans and AI in relation to one another (human and machine tacit knowledge), as well as two types of AI systems: one that goes beyond human intelligence and one that is necessarily and inherently tied to it. It finally highlights how humans and AI can augment their capabilities and intelligence through synergistic human–AI interactions (i.e., human-augmented AI and augmented human intelligence), resulting in hybrid intelligence, and concludes with a future-looking research agenda.

Keywords

Artificial intelligence, human intelligence, hybrid intelligence, human–AI interaction, human-augmented AI, augmented human intelligence

Introduction

Situated at the intersection of myth, hype, and technological reality, artificial intelligence (AI) is often used as a catchall term—a descriptor applied to a range of computational systems for strategic or marketing purposes regardless of whether or not they contain any AI at all (Musa Giuliano, 2020; Newlands, 2021; Schmelzer, 2019). For example, a survey by MMC Ventures (2019) indicated that forty percent of startups in Europe that are presented as “AI companies” did not draw upon AI “in a way that is material to their businesses.” As Jordan (2019) argues, the use of the term AI is an “intellectual wildcard” and academic, technical and political discussions have accordingly jumped on the AI bandwagon, variously adopting and adapting the term AI based on their own interests. However, ongoing misclassifications coupled with hyperbolic perceptions about what AI can and cannot do (e.g., AI systems reaching human-like intelligence in near future) can hamper legitimate efforts to derive value from AI systems across domains (Barro and Davenport, 2019). The benefits that can be derived from AI systems are growing, as has been observed in organizational settings, in public services

or in healthcare (Jarrahi et al., 2021). Yet, without a clearer or even basic understanding of the boundaries between AI and human intelligence, too many stakeholders and decision-makers may be misled by the hype around AI and miss out on the genuine benefits AI can offer. As a response to this issue, this article provides a delineation of AI and what the intelligence in AI may and may not entail.

Since the term was first coined at the Dartmouth conference in 1956, AI has undergone a profound transformation. Yet, despite these rapid changes, AI continues to encompass computing systems that exhibit some form of intelligence. One sensible way to unpack AI is, therefore, to focus on what we mean by “intelligence” within AI.

¹University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

²Nordic Centre for Internet and Society, BI Norwegian Business School, Oslo, Norway

Corresponding author:

Christoph Lutz, Nordic Centre for Internet and Society, BI Norwegian Business School, Oslo, Norway.
Email: christoph.lutz@bi.no



Integrating 70 definitions of intelligence from several disciplines, Legg and Hutter (2007) present intelligence as “an agent’s ability to achieve goals in a wide range of environments.” John McCarthy, one of the pioneers of AI, articulated AI as the ability of machines to solve problems and achieve goals in the world (McCarthy, 2012). Seen this way, AI can be understood as a system’s ability to learn and reason from experience and its environment, and to adapt to new situations. As such, rather than being deterministic and mechanistic, AI systems integrate real-world experiences to learn and adapt (West, 2018).

Artificial intelligence in relation to human intelligence

As De Cremer and Kasparov (2021) recently noted, “our principal challenge as business people is to anticipate what AI means in relationship to how humans think and act.” In exploring that relationship, we can discern three interdependent areas represented as the yellow, orange, and blue-shaded areas within Figure 1. In relation to human intelligence (yellow), we can identify two types of AI: one that goes beyond human intelligence (blue) and one that is inherently intertwined with it in a form of hybrid intelligence (orange). Through interactions between humans and AI, both can augment their capabilities and intelligence, marked by the orange overlap area in Figure 1. The main products of this synergistic relationship are human-augmented AI and augmented human intelligence.

Human intelligence

In a recent article titled “Why AI is Harder Than We Think,” Mitchell (2021), a computer scientist, argues that one of the reasons behind prevalent misunderstandings about AI and the several ebb-and-flow hype cycles in the past stems from a lack of understanding about the nature and complexity of human intelligence. The power of AI in replicating human intelligence is overhyped and associated with “high levels of over-exuberance and media attention” (Jordan, 2019). Similarly, Kerr et al. (2020) note how public expectations about AI and technological reality diverge, while a recent IBM report demonstrates a clear intention-action gap when it comes to the (ethical) application of AI in organizational settings (Goehring et al., 2022).

Put simply, humans possess “general intelligence” in being able to comprehend and analyze various situations and stimuli, to ideate, create and imagine. The intelligence projected by AI systems is predominantly task-centered (Narayanan and Kapoor, 2022). AI systems can reveal correlations in data but are not as skillful as humans in determining causation. Even though AI’s performance on

several tasks has surpassed and will continue to surpass human capacities, for AI to advance beyond narrow tasks requires holistic and contextual thinking that remains a core human capability. Humans benefit from “analogical thinking,” which enables us to envision common relational systems in different situations, thus effectively integrating prior experiences into new knowledge problems and novel domains. Similar high-level reasoning and thinking, however, remain elusive for AI systems (Jordan, 2019).

As a result, humans continue to provide hard-to-imitate emotional, social and cooperative intelligence, particularly in situations with loosely defined goals. Most AI applications in the wild (i.e., outside the computer lab) involve unstructured tasks and unknown elements. These, accordingly, may require both higher-level reasoning and intuitive decision-making, which is the unique prerogative of humans (Jarrahi, 2018).

Although intelligent machines can provide greater analytical capabilities in relation to quantitative data collection and analysis, it is humans who engage in the more effective qualitative, holistic, and intuitive analysis of data. This enables humans to imagine and anticipate. For example, humans can utilize tacit knowledge in balancing trade-offs to maximize the interests of various stakeholders in a decision-making situation. Tacit knowledge embodies personal wisdom, insight, and intuition and is intertwined with a wealth of experience over time. By definition, tacit knowledge is hard to express or extract (made explicit); as Polanyi (1966) famously noted: “we can know more than we can tell” (p. 4). Consequently, much of humans’ tacit knowledge is impossible to be replicated by AI.

Human intelligence is also embodied in the physical and biological sense, relying on the corporeality of the brain, nervous system, and complex sensory system that constitutes the human body (Dreyfus, 1992). AI systems, especially if disembodied and based on abstract representations, will fail to replicate human intelligence because they lack similar physical and biological foundations.

Artificial intelligence

AI, as represented by the blue area in Figure 1, can be qualitatively different from human intelligence or surpass it by a level of performance unattainable by humans. This is particularly so in classification tasks and in speed or reach of execution. Emerging algorithms offer forms of intelligence and problem-solving that are unmatched by humans in terms of speed and scale (e.g., fraud detection across multiple transactions). AI systems may also go beyond the bounds of existing frameworks, assumptions, and hard-coded rules formulated and provided by humans, eventually developing their own learning logic, rather than solely relying on human intelligence to accomplish tasks. Although AI may get things right and provide outputs with high efficiency, AI lacks the ability to explain how

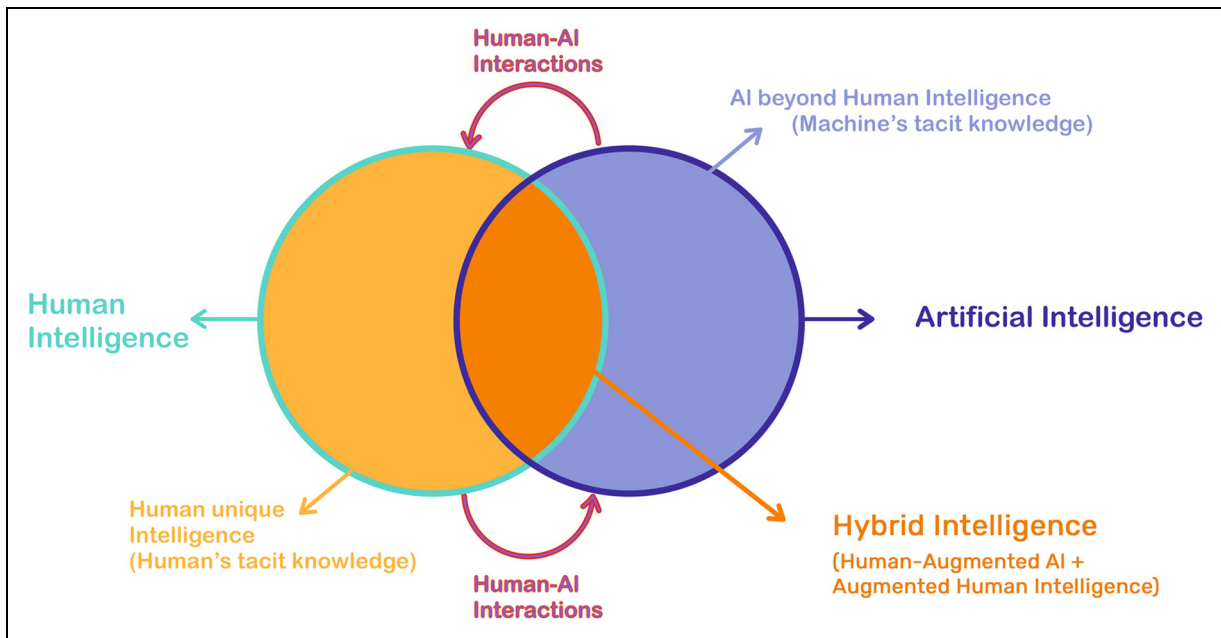


Figure 1. Artificial intelligence in relation to human intelligence.

the process works—what can be called “machine tacit knowledge.” This is the knowledge that the machine develops through self-learning but cannot necessarily codify and transfer to humans or other machines (Burrell, 2016; Felzmann et al., 2019).

Even though deep neural networks are inspired by the way neurons function in the brain, this form of intelligence is not necessarily based on how humans process information and make decisions (Hawkins, 2021). This is in line with McCarthy’s (2012) argument that AI can learn by observing and imitating humans or developing intelligence by going through methods that are not observed in humans or analytical capabilities well beyond what people can achieve. However, this form of intelligence is predominantly task-centric, meaning AI algorithms could perform extremely well in rather specific, narrow tasks with clearly externally definable goals and metrics.

Hybrid intelligence

The overlap (orange area in Figure 1) can be understood as “hybrid intelligence” (Dellermann et al., 2019), which incorporates two distinct outcomes: (1) human-augmented AI (in amplifying machine intelligence) and (2) augmented human intelligence (in amplifying human intelligence). These outcomes are fostered through continuous human–AI interactions.

Human-augmented AI. Human-augmented AI refers to AI systems that are trained by humans and continuously improve their performance based on human input. Since its inception in the 1950s, one of the primary missions of

the AI community has been to mimic human capabilities such as sensory perception, natural language processing or logical reasoning. Alan’s Turing vision of “thinking machines” is exemplified in Turing’s famous Test: “computers need to complete reasoning puzzles as well as humans in order to be considered ‘thinking’ in an autonomous manner” (West, 2018). So, the best AI systems based on this persistent vision are those that get as close as possible to the level of human intelligence. Nevertheless, AI systems so far only reflect a relatively small fraction of human intelligence, so it is important to disavow the “original sin” in the field of AI, which assumes that “minds are like computers and vice versa” (Crawford, 2021).

In more cases than AI developers and evangelists would like to recognize, machine intelligence in this context continues to be augmented by human intelligence. AI’s narrow intelligence still depends on large corpora of training data generated or processed by humans (Mitchell, 2021). Recent evidence also shows that AI vendors may keep the human labor required to train and produce AI services a secret from investors or clients for strategic purposes (Newlands, 2021). In fact, many people work tirelessly behind the scenes as a form of “human computation” (i.e., using humans as computers to perform tasks that the technical systems cannot perform alone) (Gray and Suri, 2019), including those who contribute to training AI systems without knowing about it (e.g., users training the Google search engine). As a result, rather than being fully automated systems, human-augmented AI systems are “technologies of heteromation” that crucially rely on humans as indispensable mediators (Ekbja and Nardi, 2014).

Augmented human intelligence. AI systems can also, in turn, augment human intelligence. In most cases, AI systems tend to extend or amplify human capabilities by providing support systems such as predictive analytics rather than replacing them, resulting in augmented (human) intelligence. For example, personal intelligent assistants do not make decisions for users but help broaden their cognitive bandwidth by providing useful affordances for processing, filtering, sorting, and navigating expansive information landscapes (Maedche et al., 2019). Another example comes from the game of chess. Even grandmasters can play smarter when teamed up with AI. Gay Kasparov argues that partnering with an AI system amplified his capability by allowing him to focus on strategic planning and moves while the machine took care of analytical calculations of the game (De Cremer and Kasparov, 2021; see also Tomašev et al., 2022). Augmented human intelligence is, therefore, one of the major outcomes of the overlap between human and AI since it requires human and AI to work together to enhance and elevate human intelligence (Carroll, 2021).

Human–AI interactions. The outcome of effective interactions between humans and AI is human-augmented AI and augmented human intelligence. The overlap area in Figure 1 makes it clear that both humans and AI can advance through interacting with each other. Yet, in most real-world applications of AI, achieving intelligent performance requires something beyond big data, algorithmic capabilities or computational power: most notably it requires human contributions (Jarrahi et al. 2021; Tubaro et al., 2020).

As such, AI systems are sociotechnical systems that can progress in performance only if humans and AI interact and complement one another with mutual understanding (Østerlund et al., 2021). That is, machines must better understand how humans reason and operate (AI alignment), while humans must develop a better awareness of machines' decision-making logics (AI literacy) (Jarrahi et al. 2021). An important challenge that stands in the way of developing effective human–AI symbiosis is thus divergent human and machine tacit knowledge (whose logic cannot be easily articulated and communicated to the other party). For instance, the black-box nature of neural networks can be impenetrable by users or even developers of AI systems (Burrell, 2016; Felzmann et al., 2019).

Conclusion and future research agenda

Although the concept of AI remains aspirational, technological advancements in the field have enabled computational systems to continuously progress. Putting AI systems into practice and making a real impact requires a

realistic understanding of the type of intelligence AI systems can offer and how this compares with and relies on human intelligence. Figure 1 indicates two variants of AI, one that is less tied to human intelligence and one that is reliant on human–AI interaction. We also argue that even though AI systems can be inspired by human intelligence, reaching human-level intelligence is neither a sensible nor a feasible objective.

Devising autonomous systems is almost impossible for many real-world scenarios where humans need to “stay in the loop” to maintain the sociotechnical system's versatility and adaptability in relation to new tasks and environments. Some of the greatest challenges in developing AI systems have to do with bringing contextual meaning and reasoning in relation to real-world situations (Oleinik, 2019). This requires continuous human–AI interactions, a vision beyond “automating the last mile” and superior performance in narrowly defined tasks. Finally, unrealistic expectations about what AI really entails and hyperbolic claims about its capability may yield underwhelming impacts.

Before we conclude, we point to some promising areas for future research. As the field of AI is developing further, there is a constant need to scrutinize the intelligence of the machine, what it is capable of, where it needs human inputs, and how humans and AI can augment one another. For instance, using reinforcement learning, the DeepMind's AlphaZero algorithm developed knowledge of chess by only integrating rules of chess and without any prior knowledge of common human strategies. The algorithm only built on the rules and the reward system and played itself to develop the knowledge of the game. Such AI systems open up the possibility of more surprising “machine tacit knowledge” and the outcomes of these applications may provide learning opportunities for human players. Kasparov (2018: 1087) noted: the system “prioritizes piece activity over material, preferring positions that to my eye looked risky and aggressive.” Such AI systems are unlikely to make autonomous and unsupervised decisions in many domains, but future research must continue examining the ways knowledge-generated by AI systems (though focused on narrow tasks) can be put into use for offering fresh and innovative insights to human partners.

As the overlap area in Figure 1 grows in importance, human–AI interactions and how they yield hybrid intelligence and mutual augmentation deserve more scholarly attention. More empirical and conceptual attempts are needed to advance concepts such as ‘human-AI hybrids’, ‘human-AI configurations’ or ‘human-AI symbiosis’ that help capture the sociotechnical systems emerging from human–AI interactions. An example of such theoretical developments comes from recent work in ‘human-centered AI’ (Shneiderman, 2022).

With regards to human-augmented AI, a promising research program is to trace the human labor that goes into developing, training, maintaining and dissolving AI

systems (i.e., across the whole AI life cycle). Research, utilizing ethnographic and qualitative methods, has started to investigate such human input (e.g., Miceli et al., 2020; Newlands, 2021; Posada, 2022; Tubaro et al., 2020). However, more work is needed to grasp the complex ways in which humans augment AI across contexts and applications. Studies focusing on highly digitalized forms of work (Jarrahi et al., 2021), private uses (e.g., social media and content moderation, e.g., Llansó, 2020) and domestic applications such as smart speakers (Tubaro and Casilli, 2022) show that human augmentation is both a feature and a bug of AI. Future research could systematically compare application areas and technologies, for example, embodied vs. disembodied vs. embedded AI (Glikson and Woolley, 2020), in terms of what human augmentation does to the AI and what the AI does to human augmentation (Jarrahi et al., 2023).

When it comes to augmented human intelligence, important aspects that should be investigated relate to fairness, inequality, trust, literacy and privacy. Questions such as the following should be answered: How does AI augment human intelligence in social fields such as healthcare, education, and the legal system, and how (un)fair is these tendencies for different stakeholders (Trewin et al., 2019)? Who benefits from AI that augments human intelligence (Lutz, 2019)? Who is left out or suffers due to discrimination, a lack of access or undue consequences of augmentation (Veale and Binns, 2017)? What is the right level of trust in AI systems that extend human capacities, so that both overtrust (Aroyo et al., 2021) and undertrust/aversion (Dietvorst et al., 2015) are avoided? How can AI literacy be conceptualized in a holistic way, and what are the cognitive and affective prerequisites to the implementation of augmented human intelligence (Long and Magerko, 2020)? How does augmented human intelligence come with vulnerabilities in terms of surveillance and privacy (Lutz and Newlands, 2021)?

These questions and research direction show how ethical aspects are key to understanding human-augmented AI, augmented human intelligence, and human–AI interactions. Future research on the intersection of AI and human intelligence should be interdisciplinary and adopt a socio-technical understanding. The technical aspects (e.g., represented by expertise from computer science or the tech industry), social/human aspects (e.g., represented by expertise from sociology, psychology, and communication) and ethical and legal aspects (e.g., represented by expertise from philosophy and law) should be studied in conjunction and close synergy within teams that prioritize holistic understandings of AI in context (Bailey and Barley, 2020).

Acknowledgements

The second and third authors are funded within the Research Council of Norway project “*Future Ways of Working in the Digital Economy*”

(project nr. 275347), and the second author is also funded within the Research Council of Norway project “*Algorithmic Accountability: Designing Governance for Responsible Digital Transformations*” (project nr. 299178).

Acknowledgments

We appreciate Gary Marchionini, Ronald Bergquist and Min Kyung Lee for their feedback on earlier drafts of this article. We are also grateful to Zhaleh Ghalebeygi, Samira Momenipour and Reza Bagherzadeh for their help with Figure 1.

Declaration of conflicting interests


The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Norges Forskningsråd (grant number 275347, 299178).

ORCID iDs

Mohammad Hossein Jarrahi  <https://orcid.org/0000-0002-5685-7156>

Christoph Lutz  <https://orcid.org/0000-0003-4389-6006>

References

- Aroyo AM, De Bruyne J, Dheu O, et al. (2021) Overtrusting robots: Setting a research agenda to mitigate overtrust in automation. *Paladyn Journal of Behavioral Robotics* 12(1): 423–436.
- Bailey DE and Barley SR (2020) Beyond design and use: How scholars should study intelligent technologies. *Information and Organization* 30(2): 100286.
- Barro S and Davenport TH (2019) People and machines: Partners in innovation. *MIT Sloan Management Review* 60(4): 22–28.
- Burrell J (2016) How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society* 3(1): 1–12.
- Carroll N (2021) Augmented intelligence: An actor-network theory perspective. In European Conference on Information Systems (ECIS), paper nr. 1219. Available at: https://aisel.aisnet.org/ecis2021_rp/37/ (accessed 7 November 2022).
- Crawford K (2021) *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, CT: Yale University Press.
- De Cremer D and Kasparov G (2021) AI should augment human intelligence, not replace it. *Harvard Business Review*, 18 March. Available at: <https://hbr.org/2021/03/ai-should-augment-human-intelligence-not-replace-it> (accessed 7 November 2022).
- Dellermann D, Ebel P, Söllner M, et al. (2019) Hybrid intelligence. *Business & Information Systems Engineering* 61(5): 637–643.
- Dietvorst BJ, Simmons JP and Massey C (2015) Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General* 144(1): 114.

- Dreyfus HL (1992) [1972] *What Computers Still Can't Do: A Critique of Artificial Reason*. Cambridge, MA: MIT Press.
- Ekbia H and Nardi B (2014) Heteromation and its (dis)contents: The invisible division of labor between humans and machines. *First Monday* 19(6). Available at: <https://journals.uic.edu/ojs/index.php/fm/article/view/5331/4090> (accessed 7 November 2022).
- Felzmann H, Fosch Villaronga E, Lutz C, et al. (2019) Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society* 6(1): 1–14.
- Glikson E and Woolley AW (2020) Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals* 14(2): 627–660.
- Goehring B, Rossi F and Rudden B (2022) AI ethics in action. IBM Institute for Business Value, 31 March 2022. Available at: <https://www.ibm.com/downloads/cas/4DPJK92W> (accessed 7 November 2022).
- Gray ML and Suri S (2019) *Ghost Work: How to Stop Silicon Valley from Building a New Global Underclass*. Boston, MA: Houghton Mifflin Harcourt.
- Hawkins J (2021) We'll never have true AI without first understanding the brain. MIT Technology Review, 3 March 2021. Available at: <https://www.technologyreview.com/2021/03/03/1020247/artificial-intelligence-brain-neuroscience-jeff-hawkins/> (accessed 7 November 2022).
- Jarrahi MH (2018) Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons* 61(4): 577–586.
- Jarrahi MH, Lutz C, Boyd K, et al. (2023) AI in the work context. *Journal of the Association for Information Science and Technology*. forthcoming.
- Jarrahi MH, Newlands G, Lee MK, et al. (2021) Algorithmic management in a work context. *Big Data & Society* 8(2): 1–14.
- Jordan MI (2019) Artificial intelligence—the revolution hasn't happened yet. Harvard Data Science Review, 2 July. Available at: <https://hdsr.mitpress.mit.edu/pub/wot7mkc1/release/5> (accessed 7 November 2022).
- Kasparov G (2018) Chess, a drosophila of reasoning. *Science (New York, N.Y.)* 362(6419): 1087.
- Kerr A, Barry M and Kelleher JD (2020) Expectations of artificial intelligence and the performativity of ethics: Implications for communication governance. *Big Data & Society* 7(1): 1–12.
- Legg S and Hutter M (2007) A collection of definitions of intelligence. In: Goertzel B and Wang P (eds) *Advances in General Artificial Intelligence: Concepts, Architectures and Algorithms and Applications*. Amsterdam: IOS Press, pp.17–24.
- Llansó EJ (2020) No amount of “AI” in content moderation will solve filtering's prior-restraint problem. *Big Data & Society* 7(1): 1–6.
- Long D and Magerko B (2020) What is AI literacy? Competencies and design considerations. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp.1–16.
- Lutz C (2019) Digital inequalities in the age of artificial intelligence and big data. *Human Behavior and Emerging Technologies* 1(2): 141–148.
- Lutz C and Newlands G (2021) Privacy and smart speakers: A multi-dimensional approach. *The Information Society* 37(3): 147–162.
- McCarthy J (2012) What is AI? Basic Questions. Available at: <http://jmc.stanford.edu/artificial-intelligence/what-is-ai/index.html> (accessed 7 November 2022).
- Maedche A, Legner C, Benlian A, et al. (2019) AI-based digital assistants. *Business & Information Systems Engineering* 61(4): 535–544.
- Miceli M, Schuessler M and Yang T (2020) Between subjectivity and imposition: Power dynamics in data annotation for computer vision. *Proceedings of the ACM on Human-Computer Interaction* 4(CSCW2): 1–25.
- Mitchell M (2021) Why AI is harder than we think. arXiv preprint, arXiv:2104.12871. Available at: <https://arxiv.org/abs/2104.12871> (accessed 7 November 2022).
- MMC Ventures (2019) The state of AI 2019. MMC Ventures. Available at: <https://www.stateofai2019.com/> (accessed 7 November 2022).
- Musa Giuliano R (2020) Echoes of myth and magic in the language of artificial intelligence. *AI & Society* 35(4): 1009–1024.
- Narayanan A and Kapoor S (2022) Why are deep learning technologists so overconfident? AI Snake Oil. Available at: <https://aisnakeoil.substack.com/p/why-are-deep-learning-technologists> (accessed 7 November 2022).
- Newlands G (2021) Lifting the curtain: Strategic visibility of human labour in AI-as-a-service. *Big Data & Society* 8(1): 1–14.
- Oleinik A (2019) What are neural networks not good at? On artificial creativity. *Big Data & Society* 6(1): 1–13.
- Østerlund C, Jarrahi MH, Willis M, et al. (2021) Artificial intelligence and the world of work, a co-constitutive relationship. *Journal of the Association for Information Science and Technology* 72(1): 128–135.
- Polanyi M (1966) *The Tacit Dimension*. Chicago, IL: University of Chicago Press.
- Posada J (2022) Embedded reproduction in platform data work. *Information, Communication & Society* 25(6): 816–834.
- Schmelzer R (2019) Is machine learning really AI part 2. Forbes, 29 November. Available at: <https://www.forbes.com/sites/cognitiveworld/2019/11/29/is-machine-learning-really-ai-part-2/> (accessed 7 November 2022).
- Shneiderman B (2022) *Human-centered AI*. Oxford: Oxford University Press.
- Tomašev N, Paquet U, Hassabis D, et al. (2022) Reimagining chess with AlphaZero. *Communications of the ACM* 65(2): 60–66.
- Trewin S, Basson S, Muller M, et al. (2019) Considerations for AI fairness for people with disabilities. *AI Matters* 5(3): 40–63.
- Tubaro P and Casilli A (2022) Human listeners and virtual assistants: Privacy and labor arbitrage in the production of smart technologies. In: Graham M and Ferrari F (eds) *Digital Work in the Planetary Market*. Cambridge, MA: MIT Press, pp.175–190.
- Tubaro P, Casilli AA and Coville M (2020) The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society* 7(1): 1–12.
- Veale M and Binns R (2017) Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society* 4(2): 1–17.
- West DM (2018) What is artificial intelligence? Brookings. Available at: <https://www.brookings.edu/research/what-is-artificial-intelligence/> (accessed 7 November 2022).