



# AI and the Future of Work Workshop Report

4 & 5 December 2023

Lily Rodel, Ekaterina Hertog, Charlotte Unruh, Daniel Susskind<sup>1</sup>

<sup>1</sup> This report was written by Lily Rodel with input from Ekaterina Hertog. Ekaterina Hertog, Daniel Susskind, and Charlotte Unruh conceptualized and organised the workshop on which this report is based. All authors read and commented on the final version of this report. We thank all participants for their input and contribution to an insightful event.



**Dr Ekaterina Hertog**

Associate Professor in AI and Society, Institute for Ethics in AI and Oxford Internet Institute, University of Oxford



**Dr Charlotte Unruh**

Early Career Research Fellow, Institute for Ethics in AI, University of Oxford. In January 2024, Charlotte will take up a position as Lecturer in Philosophy at the University of Southampton.



**Professor Daniel Susskind**

Professor Daniel Susskind, Research Professor in Economics, King's College London, and a Senior Research Associate, Institute for Ethics in AI at Oxford University.

	<b>Page</b>
<b>Executive Summary</b>	3
<b>Introduction</b>	5
<b>Background</b>	6
<b>Workshop Themes</b>	7
- <b>Meaningful work</b>	7
- <b>Justice</b>	8
- <b>Governance</b>	9
- <b>When AI is not ‘just work’</b>	11
<b>Future directions</b>	13
<b>Acknowledgements</b>	14
<b>Works cited</b>	15

© 2023 Rodel, L., Hertog, E., Unruh, C., Susskind, D. All rights reserved

**Keywords:** AI, work, the future of work, meaningful work, care work, governance

**Citation:** Rodel, L., Hertog, E., Unruh, C., Susskind, D. (2024) AI and the Future of Work Workshop Report. Institute for Ethics in AI, University of Oxford

The Institute for Ethics in AI, University of Oxford, hosted a two-day workshop on the topic of AI and the future of work. This workshop discussed many different facets of paid work: the white-collar workplace, the AI supply chain, meaningful work vs menial work, the creative industries, and the care sector amongst others. These different contexts implicate different actors, dynamics, policies, and indeed challenges when intersected with AI implementation. During this workshop, work and its role and meaning in social life were placed under the microscope, with participants sharing diverse understandings: to some, work is something you would gladly pay another to do, whilst to others, work is a source of meaning, identity, and a source of collective action. For many, work is both. This workshop shed light on how an understanding of the role of AI in work can open up conversations around what work is and, in a technologically enabled future, what work could and should be like. These discussions highlight the role of academics, policymakers, and business leaders in the face of rapid technological advancement to help societies in understanding the choices they face, and to illuminate the possible futures ahead.

We summarise our key discussion points and insights below under four key themes.

**Meaningful work:** The emergence and rapid adoption of generative AI in the workplace has called into question the role of work and workers. Though many have warned of automation and its impact on the labour market, several participants felt that instead workplace augmentation may become more widespread.

Augmentation may replace drudgery and repetition at work and create more time for workers to spend on tasks they find ‘meaningful’, it may also contribute to increased work intensification and shifting expectations of what is possible for workers to accomplish. Further, the implementation of AI augmentation in the workplace reveals a ‘jagged frontier’ where different configurations of workers, tasks, and AI placement generates distinct results. In the future, participants urged business leaders to prioritise long-term growth rather than short-term automation gains when it comes to AI implementation in the workplace. Others recommended further empirical research on the daily realities and expectations of workers in AI-enabled workplaces.

**Justice:** AI development and implementation often follows business interests that do not necessarily intersect with ethical considerations of justice and obscures the vital role of workers, both in white-collar workplaces in advanced economies and in the global AI supply chain.

A spotlight on justice adds nuance to the conversation on meaningful work: meaningful work for who? For some AI workers, such as those working in so-called ‘AI sweatshops’, requests for fair pay, a safe working environment, and a formalised working contract are more pressing than requests for meaningful work. This led some workshop participants to advocate for democratising workplaces in general and of AI development and production process more specifically. This approach, it was argued, would enable the creation of more ethical AI.

**Governance:** Governance should focus on the base standards that regulations can introduce, such as the fundamental rights of privacy and non-discrimination. AI governance processes should be made transparent to workers. However, the introduction of AI to the workplace has generated complicated questions around where the responsibility for regulation should lie.

Due to the rapid rate of change, organisations and institutions who build and procure these technologies are under increasing pressure to ameliorate potential risk. Participants identified two levels of governance: the local layer, referring to those who deploy AI, and the global layer, referring to macro issues such as workforce, democracy, and existential questions. Though participants agreed that regulation should adhere to already enshrined rights, the question of who should have the right to design how AI is used has become a central concern, next to the issue of who should assume the responsibility to enforce the regulation: the state, the firm, the workers themselves?

**When AI is not ‘just work’:** This section focused on cases where AI is adopted in industries that rely heavily on human connection (like health care, education, or social care). Whilst some participants warned of the risks of increased datafication in care sectors, others predicted that AI could expand the capability of the sector and lead to greater human connection. Ultimately, technology is not a panacea – the way AI is positioned in these sectors will determine its impact and should always be based on consultations with the people providing and receiving care, i.e. both workers and users.

The discussions around this topic largely focused on the paid care sector, including healthcare and social care, but the speakers highlighted the parallels and connections with unpaid care provided by families. In an ageing population, care needs are increasing, but within the current system in the UK, as well as in many other Global North countries care work is understaffed, undervalued, and underfunded. Adoption of AI that focuses on cost-cutting in this sector continues the current disregard for care work and could lead to an ‘accountability gap’<sup>1</sup>. Rather than replace humans, in this sector AI should focus on helping carers to do ‘what they do best’. Automating some parts of care could create more capacity within the system for greater levels of human connection, but this will not happen automatically. Ultimately, the impact of AI on the care sector depends on the aims and understandings of decision-makers. Academics and policymakers, therefore, must focus on gathering data and empirical evidence to support greater human connection in the care sector.

---

<sup>1</sup> Workshop participants used the term ‘accountability gap’ to refer to the issues in implementing unregulated AI within the care sector that result from fuzzy definitions of who should be held accountable should there be any negative outcomes for users, patients, and/or clients.

This report encapsulates the rich discussions and diverse perspectives that emerged during a two-day workshop hosted by the Institute for Ethics in AI, University of Oxford, on the topic of AI and the future of work. This workshop brought together leading academics from different disciplines as well as high-profile policy and business professionals to discuss the ethical issues that arise in increasingly data-driven workplaces.

This workshop focused on showcasing how foundational research in philosophy and related areas can be brought to bear on questions of practical importance on the topic of AI and the future of work. Discussions revolved around the following questions:

- What is the normative significance of work?
- What do we stand to lose, or gain, in a world where a lot of work is outsourced to AI?
- How does the use of technologies change the nature of work in standard and non-standard employment?
- What are the challenges that generative AI brings to the workplace?
- How can these changes be governed responsibly?

Following Chatham House rules, the following report takes stock of the presentations and discussions throughout the workshop, as well as the future-facing questions raised by our diverse group of participants. While we do not name them, participants hailed from academia and the private and public sector, working locally in the UK as well as in international contexts.

This report is structured as follows: we begin by providing a brief background to the topic of AI and work, followed by a discussion of the four key themes discussed during the workshop: meaningful work, justice, governance, and when AI is not ‘just’ work. We close with open questions and calls for future research and action.

In an era dominated by transformative artificial intelligence (AI) technologies, the landscape of work is undergoing a profound shift, prompting critical reflections on its implications for organisations, regulators, and workers. Though AI is not new, the launch of OpenAI's large language model (LLM) ChatGPT 3.5 in November 2022 (OpenAI 2022) brought generative AI to the forefront due to its ground-breaking ability to perform tasks once thought to be exclusive to human capabilities. Other Generative AI technologies include text-to-image models such as Dall-E that can produce unique photorealistic images from text and image prompts. These Generative AI models are significant for their accessibility: as natural language processors, they perform as if they understand text and spoken words in a similar way to human beings and, therefore, do not require any specialist software knowledge to use.

AI's capabilities have sparked vigorous debates on the future of work and the fate of workers. As we grapple with the integration of AI into various sectors, and the excitement at the possibility of Generative AI for work reaches a fever pitch, questions and concerns emerge regarding the potential replacement of human workers through technological automation, or the possibilities for augmentation and human-in-the-loop capabilities in the workplace. Tech-enabled smart workplaces, while promising increased efficiency, also bring concerns of surveillance, intensification, amplified expectations, and degraded work conditions. The implementation and governance of AI takes place on uneven terrain, and this workshop highlighted the need for perspectives on AI that address power asymmetries on a global level, in addition to ethical issues arising at a local and firm level.

Workshop discussions revolved around four key themes, stated below:

**Meaningful work:** The impact of AI on work is explored through debates on whether automation will replace or enhance tasks, raising essential normative questions about the value and role of work in individuals' lives, and who can and cannot access 'meaningful' work.

**Justice:** Participants scrutinised the power dynamics within AI implementation, emphasising a democratic justice frame to guide responsible innovation and design and urging a reconsideration of the political role of tech firms in producing global inequalities, and of their responsibility in addressing them.

**Governance:** Participants underscored the need for effective governance of AI integration into workplaces, with discussions delving into local and global levels, complexities around the individualisation of risk and responsibility, and the democratisation of decision-making to address the challenges of responsible AI use.

**When AI is not 'just' work:** Participants delved into the nuanced role of AI in jobs requiring human connection or creativity, emphasising the importance of balancing potential benefits with risks, and highlighting the necessity of involving carers and those cared for in AI design decisions to address accountability gaps and preserve human connections.

## Meaningful Work

Considering the role, implementation, and future of AI in work sheds light on several core issues and questions that transcend AI itself: the meaning of work and labour, the value of work, the divisions of labour across local, global, and organisational scales, who makes decisions that shape the daily lives of workers, and where is the risk and responsibility for these decisions concentrated? This workshop revealed that AI is an amplifier that can reveal and lead us to question taken-for-granted assumptions about work and the workplace. These discussions highlight the role of academics, policymakers, and business leaders in the face of rapid technological advancement to help societies in understanding the choices they face and illuminating the possible futures ahead.

AI is increasingly changing the way we work. Techno-utopians celebrate and AI doomers warn us of the transformative capabilities of emergent technologies, which enact change at an unprecedented rate. In recent months, the emergence and rapid adoption of Generative AI in the workplace has expanded common conceptions of the technological remit. AI can now compose music, make medical decisions, and create art. In light of these sociotechnical changes, workshop participants discussed: what does this mean for work and workers? Will AI replace workers, and what effect might this have on the division of labour? Workshop participants addressed the contemporary fears around AI and its impact on the workforce, debating whether AI could replace work altogether, or whether it might augment workers. At their core, these discussions were concerned with essential normative and ethical questions to do with work: what is work? What is the meaning of work in the lives of workers and in society itself? What could and what should it be?

Will AI replace workers, or will it augment their work? During the workshop, participants weighed up the assumptions at the heart of the automation vs augmentation debate. Fundamentally, distinct imaginations of a utopian future shaped by normative understandings of the role of work characterise the different positions in the debate. Participants examined the prevalent view amongst economists that work is something that the rational human would prefer not to do and thus may benefit from automation. Radical post-work views espoused by thinkers such as David Graeber deride ‘bullshit jobs’; if work is a ‘bad’ thing, something many would avoid, then automation could bring the opportunity to escape alienating work.

Others complicated these perspectives with recent literature on so-called ‘meaningful work’ and its role in identity, belonging, and the subjective human experience. If certain aspects of or tasks involved in a particular job are viewed as meaningful<sup>2</sup> and others as a kind of Marxian alienated drudgery, then technology could

---

<sup>2</sup>Participants noted the importance of situating the ‘meaningfulness’ of work within both the individual experience – referring to the subjective attachment to objectively valuable people and things (George and Park 2016), made legible through culturally-



contribute to the flourishing of non-alienating work in the future by allowing workers to focus on tasks they found meaningful, creative, or innovative, thus augmenting the capability of workers. For instance, workshop contributors suggested that the replacement of repetitive and time-consuming tasks in the creative industries by AI may not lead to job loss, but instead open up more time for innovation and ideation from human workers. Notably, even with this optimistic vision of the role AI can play in the future of work there are associated risks, such as work intensification— if AI replaces human workers in tasks that are routine and boring, but easy, and leaves humans only with interesting, but challenging parts of their jobs. If management does not pay attention to the AI changes in the nature of work, rather than improving the quality of work experience, AI may lead to worker burnout.

The results of implementing human-in-the-loop or symbiotic relationships with AI in the workplace are not straight-forward – participants highlighted research on the ‘jagged frontier’ of AI (Dell’Acqua et al 2023) and the distinction between high discretion and low discretion automation (Hopf et al 2022), which demonstrate the uneven terrain of AI at work. Further, questions around the division of labour are crucial: who gets to do the meaningful work? Whose role will be augmented and whose replaced? The replacement of tasks by AI may lead to developments such as a Universal Basic Income (UBI) or a four-day working week, or they may amplify expectations of productivity in lead to the intensification of work. Participants highlighted the need for business decision-makers to think long-term when deciding how to invest in AI in the workplace; short-term automation gains miss opportunities for growth, innovation, and re-invention. Further empirical research on the daily realities and expectations of workers in AI-enabled workplaces is required to understand the human impact of the automation vs augmentation debate.

## Justice

Following on the theme of meaningful work, workshop participants also examined the implications of AI for justice, fairness, and equity in and of the workplace. An insight into AI at work offers a window into the tangled and multi-scale power relations upon which technology is built and implemented. A justice frame can, therefore, be used to inspire responsible AI innovation and work. Participants also considered what role work might play in the cultivation of justice<sup>3</sup> and what might be lost if jobs and workplaces are replaced by automation. One of the key questions of this workshop was: can AI make work better? Participants questioned: better for who? Whose perspectives and objectives are privileged in the discourse around AI at work?

---

located meaning-making frameworks - as well as with collective and public judgements and contributions of work towards the ethics of autonomy, freedom, and dignity.

<sup>3</sup> Utilizing the work of Rawls and others, participants suggested that work is fundamental as an arena for the cultivation of justice, tolerance, and co-operation. This has further implications for the possibility of UBI opened up by previous discussions on meaningful work, and instead may suggest that a Conditional Basic Income would lead to a more ‘just’ society.

The ethics of AI production and design often follow business interests which do not necessarily intersect with democratic justice or ‘the public good’. For example, AI that improves organisational productivity through worker monitoring and surveillance may have negative consequences for worker wellbeing. Questions around justice can thus have practical relevance and guide the creation of technology. Rather than discussing technology as if the direction of development takes a predetermined course of instrumental rationality, responsible design can circumvent the risk that AI worsens work conditions through responsible innovation that addresses workers not management. Participants called on the work of thinkers such as Castoriadis who have examined how technology might be used for worker power, connecting his work with empirical examples of existing worker-led collective movements such as the creation of the Dynamo platform as a site for mobilising Amazon Mechanical Turk workers. Considerations of justice can thus demystify the inherent power relations that produce and inform the direction of AI and help to guide further AI design.

An insight into ‘better for who’ can also shine a spotlight onto the often-hidden labour that is fundamental to AI: the data labelling and content moderation outsourced to emerging and low-income economies in the Global South. Workshop participants highlighted the importance of expanding the frame of the discussion of AI at work beyond white-collar workplaces to include an acknowledgement of the global post-colonial supply chain that AI is built on. For workers in AI sweatshops in the Global South, working conditions are so poor that conversations around meaningful work and augmentation are much less relevant than essential requests for fair pay, a safe working environment, and a formalised worker contract.

The stark inequality of workers on an organisational, national, and global scale forms the basis on which AI and other technology is built – some workshop participants urged us to reconsider the increasingly powerful role of tech firms play in society as political actors. Indeed, what these discussions imply is the importance of considering work as a fundamentally political experience; a site for self-expression, for meaning-making, and also the site of economic exploitation and potential alienation. A focus on democratic justice at work encourages analysts to understand the firm as a political entity and work as a political experience – workers are governed by the rules and decisions of the firm, but at the same time are the labour investors that enable the survival of the firm<sup>4</sup>. Is the answer the democratisation of the firm and of capitalism itself?

## **Governance**

How can and how should AI development and its use at the workplace be governed? Is there a need for a new governance of AI at work? The rapid rate at which AI and other emergent technologies are being developed and incorporated into businesses puts the people, organisations, and institutions building and

---

<sup>4</sup> For a deeper exploration of this line of thinking see Ferreras Isabelle, 2017, *Firms as Political Entities*. Cambridge University Press

procuring these technologies under pressure to take steps that mitigate potential risk. Due to the speed of change and the diversity of contexts within which AI might be applied, participants agreed that the focus of governance should be on the base standards that regulations can introduce – such as the fundamental rights of privacy and non-discrimination which are already enshrined in regulation, and on adequate procedural norms.

Firstly, workshop participants identified two levels of governance: the local and the global. These levels contain different requirements and implications for the role of governance and the relevant actors. The local level refers to what governance might look like for those who deploy AI. Local governance can be further split into three layers: (1) the leadership and governance framework, which refers to a set of principles that translate into policies that allow them to enforce those principles; (2) the use and execution layer, which refers to how ‘fair’ use might be defined and executed; and finally (3) the enterprise layer, which refers to where organisations might be using AI and could include an inventory of where and which AI is used. Workers are centrally concerned at these three layers, and the pressing issue of their (absence of) rights to bear on these developments was highlighted. Hence, these governance layers must be supported on a local level by workforce education and training.

However, in practice, issues around local governance are more complex. If decisions in an organisation are being automated, where does the responsibility for regulation lie? Several organisations have tackled this by forbidding the use of large language models (LLMs), like Chat GPT, yet research suggests that many workers are using these tools without disclosing this to their management. If work tasks are being augmented by LLMs, who is incurring the risk of possible misinformation, challenges around intellectual property, or hallucinations (i.e. plain errors), and who is reaping the benefit of potentially improved productivity and efficiency? Participants observed the increasing tendency for companies to move towards AI governance processes due to fear of reputational damage if usage goes wrong, rather than due to a concern with the deployment of AI ‘for good’.

Other organisations circumvent issues to do with governance and regulation by defining themselves as ‘platforms’, thus moving away from formal employment and shifting the burden of risk and responsibility to individual platform workers. For example, digital care platforms exclude their workers from social rights and protection, leading to greater socio-economic risks for the worker such as increased unpaid labour, increased exposure to the risk of scamming and sexual harassment, and increased unpaid socio-emotional labour due to platform affordances that feature client ratings and reviews, but do not allow the workers to rate and leave feedback on clients. Further, the specific experiences of workers, whether in formalised or informalised employment, often intersects with specific macro-regulatory context – the so-called ‘global’ layer of governance.

The global layer of governance refers to issues around AI governance that organisations currently do not address: issues around workforce, democracy, and existential threats that AI might pose or accelerate. Workshop participants highlighted the need to strengthen the regulatory role of the state, democratise the firm power structure internally, and enhance worker agency. Fundamentally, not only is an understanding of how and where to introduce regulation important, but addressing who should have a say in regulation becomes central. The discussions on this topic all reflected the increased need for workers to understand not only how AI might make their work more efficient, but how it is used in their jobs such as in hiring, firing, or promotion, in the allocation of tasks, and in security decisions.

### **When AI is not ‘just work’**

Can and does AI have a role to play in jobs that rely heavily on human connection? What role should this be, and what are the challenges? Developments in the capabilities of AI have led many to imagine a future where machines might provide care and company to the young, elderly, vulnerable, or sick. Workshop participants shared contrasting perspectives on these deeply ethical issues: whilst some highlighted the potential dangers of increasing technological dependence and datafication in care sectors, others suggested that utilising AI within care sectors could expand the sectors’ capabilities, reduce the workload of its overburdened workers, and lead to an increase in opportunities for human connection. Workshop participants agreed that technology is not a panacea. Ultimately, participants called attention to the importance of the way AI is positioned within these sectors – similarly to previous discussions on meaningful work, does AI replace or does it augment care work? What is the ultimate goal behind the adoption of AI? The quality of care tends to suffer when the core focus is on cost-cutting. Any use of AI in these settings should be based on consultation with both the people who receive and provide care.

In ageing populations, care needs increase. Yet within the current economic system, care work is understaffed, undervalued, and underfunded, leading policymakers and academics to consider the role that AI and emergent technologies might play in filling the care needs gap. However, workshop participants warned against the framing of technology as the ultimate solution for the deep-rooted problems within the care sector. Workshop participants suggested that if AI implementation within the care sector builds on the current disregard for the importance of the care sector, AI adoption would risk resulting in an ‘accountability gap’ within the sector. Linking to the earlier discussion on governance and highlighting the associated issues around risk and responsibility, the workshop participants noted that the current trends governing AI adoption in care in the UK seems to be informed by the desire to avoid tackling structural issues. To address these concerns, participants recommended prioritising the views and experience of those being cared for and their care workers in the design and use of AI in the care sector. For example, children in the healthcare system are being increasingly datafied, with no real option to opt-out, and this is happening against the background of very little research on how children and their families feel about their data being used in a public sector context.

Workshop participants also discussed the role of human emotions and connection –which many consider irreplaceable by technology. This discussion tapped back into the distinction made between augmentation and replacement. For example, in the education sector, participants felt that AI should be used to ‘help teachers do what they do best’, rather than duplicating what should be considered uniquely human-facing tasks. In the healthcare sector, workshop participants debated the potentiality of AI augmentation for creating new opportunities for human connection. For example, the increasing move to virtual wards within the NHS can enable more connection between families and those being cared for, and the use of AI by doctors could reduce their workload to increase the time they are able to spend with patients. Fundamentally, the impact of AI on opportunities for human connection hinges on the logic and rationality structuring its use – will AI be used primarily to improve the bottom line, increase productivity, and efficiency, as captured through quantitative measurements? Or will it be used to open up new capabilities for human flourishing within struggling systems?

To conclude this report, we review prompts and questions raised or left unanswered through our workshop to provide guidance for future research.

- What kind of a technologically-enabled future do we want? Whose vision of the future is most prominent, and how can we imagine alternative futures? How are we making sure that service-providers (workers, school teachers, caregivers, ...) and service-users (customers, patients, children, ...) have the right to shape these futures, and provide adequate feedback on the AI reality that is already in place?
- How can philosophers, social scientists, business leaders, and technology experts work together in a multi-disciplinary way to address concerns to do with the design, implementation, and governance of AI in the workplace?
- We need more empirical research into the daily experiences of workers in AI-enabled workplaces is required to understand the shifting normative expectations of work amidst technological change.
- We need more evidence-based, independent assessment of the impacts and outcomes of AI on unregulated care sectors (e.g. EdTech) is needed.
- We need more empirical research into how AI impacts labour investors in the South, from informal workers to platform workers across the so-called “global value-chain” that are uniquely needed to clean data sets and produce AI but that are most likely to be exploited economically and exposed to various forms of harms, in particular emotional ones.
- We need more empirical research into how AI might impact unpaid labour – is the augmentation vs automation debate still relevant, or are new explanatory frameworks required?
- AI implementation should be informed by not only the concerns of business decision-makers or designers in Silicon Valley, but by the perspectives of the people directly affected. More ethnographic and qualitative insights into the experience of AI-enabled settings are needed: for example, how do children feel about the increased datafication of public sector services?

We extend our heartfelt thanks to all those who played a pivotal role in making this workshop a success. Special thanks to Isabelle Ferreras, who contributed to event conceptualisation, played an integral role in identifying and inviting many of the speakers, and was the keynote speaker at the event.

Our gratitude also goes to the Institute for Ethics in AI, whose extensive organisational support was invaluable. We are particularly grateful to Marie Watson, Lauren Czerniawska, Rhiannon Jones, and Paul Evans from the Institute's team for their tireless efforts and dedication behind the scenes.

This workshop and report have been funded by the Institute for Ethics in AI courtesy of the McGovern Foundation. We wish to acknowledge their generous financial support, which was crucial in bringing our vision to life.

Lastly, our sincere appreciation goes to Wadham College for generously hosting the event, and providing a welcoming and conducive environment for learning and collaboration.

Dell'Acqua et al. (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality Harvard Business School Technology & Operations Mgt. Unit Working Paper No. 24-013, Available at SSRN: <https://ssrn.com/abstract=4573321> or <http://dx.doi.org/10.2139/ssrn.4573321>

Ferreras Isabelle, 2017, Firms as Political Entities. Cambridge University Press

George, L. S., & Park, C. L. (2016). Meaning in Life as Comprehension, Purpose, and Mattering: Toward Integration and New Research Questions. *Review of General Psychology*, 20(3), 205-220. <https://doi-org.ezproxy-prd.bodleian.ox.ac.uk/10.1037/gpr0000077>

Hopf, K., Müller, O., Shollo, A., & Theiss, T. (2023). Organisational Implementation of AI: Craft and Mechanical Work. *California Management Review*, 66(1), 23-47. <https://doi.org/10.1177/00081256231197445>

OpenAI. (2022) 'Introducing ChatGPT'. OpenAI. November 30, 2022. Available at: <https://openai.com/blog/chatgpt>. (Accessed 9 December 2022).