ORIGINAL PAPER



The Ethical Implications of Artificial Intelligence (AI) For Meaningful Work

Sarah Bankins¹ • Paul Formosa²

Received: 17 February 2021 / Accepted: 25 January 2023 © The Author(s) 2023, corrected publication 2023

Abstract

The increasing workplace use of artificially intelligent (AI) technologies has implications for the experience of meaningful human work. Meaningful work refers to the perception that one's work has worth, significance, or a higher purpose. The development and organisational deployment of AI is accelerating, but the ways in which this will support or diminish opportunities for meaningful work and the ethical implications of these changes remain under-explored. This conceptual paper is positioned at the intersection of the meaningful work and ethical AI literatures and offers a detailed assessment of the ways in which the deployment of AI can enhance or diminish employees' experiences of meaningful work. We first outline the nature of meaningful work and draw on philosophical and business ethics accounts to establish its ethical importance. We then explore the impacts of three paths of AI deployment (replacing some tasks, 'tending the machine', and amplifying human skills) across five dimensions constituting a holistic account of meaningful work, and finally assess the ethical implications. In doing so we help to contextualise the meaningful work literature for the era of AI, extend the ethical AI literature into the workplace, and conclude with a range of practical implications and future research directions.

Keywords Meaningful work · Artificial intelligence (AI) · Ethical AI · Future of work · Technology and work

Introduction

Increasing organisational use of artificially intelligent (AI) technologies will influence how people experience work (World Economic Forum [WEF], 2018), including how and whether they experience meaningfulness in their work. AI is the ability of computers and other artificial entities to do things typically classified as requiring intelligence were a human to do them, such as reason, plan, problem solve, and learn from experience (Wang, 2019). Meaningful work is the perception that one's work has worth, significance, or a higher purpose (Michaelson et al., 2014), and this typically

✓ Sarah Bankins sarah.bankins@mq.edu.auPaul Formosa paul.formosa@mq.edu.au

Published online: 11 February 2023

- Department of Management, Macquarie Business School, Macquarie University, North Ryde Campus, Sydney, NSW 2109, Australia
- Department of Philosophy, Faculty of Arts, Macquarie University, North Ryde Campus, Sydney, NSW 2109, Australia

requires the coordinated exercise of varied and complex skills to benefit others. Providing opportunities for meaningful work supports positive outcomes for workers (Allan et al., 2019) and is ethically important as a basis for human wellbeing and flourishing (Bailey et al., 2019; Lysova et al., 2019). However, despite becoming an increasingly prevalent feature of workplaces, there remains a poor understanding of how AI use will influence opportunities for meaningful work and the ethical implications of such changes.

Historically technological advancements have, since at least the first industrial revolution, significantly changed opportunities for meaningful work by altering what workers do, the nature of their skills, and their feelings of alienation from or integration with the production process (Vallor, 2015). AI use will likely extend such changes, but its unique features and uses also generate new and conflicting implications for meaningful work. Optimistic accounts suggest that AI will expand the range of meaningful higher-order human work tasks (WEF, 2018), whereas more pessimistic accounts suggest that AI will degrade and even eliminate human work (Frey & Osborne, 2017). These ongoing tensions point to a lack of conceptual clarity regarding the



impacts of AI on meaningful work, leading to calls for more research in this area (Parker & Grote, 2022).

This conceptual paper aims to help address such gaps by examining how workplace use of AI has the potential to both enhance and diminish experiences of meaningful work, depending largely on the implementation choices of employers. This research is positioned at the intersection of the meaningful work and ethical AI literatures and makes two key contributions. First, we contextualise the meaningful work literature for the era of AI by developing conceptual resources to examine how the implementation of such technologies affects workers' opportunities for meaningful work and connect this assessment to the ethical implications of these changes. Second, we help remedy a neglected aspect of the ethical AI literature by offering a detailed examination of AI's implications for meaningful work.

We begin by outlining the nature of meaningful work and its ethical importance, integrating philosophical and business ethics accounts. We then examine the impacts of three paths of AI deployment—replacing some simple and complex tasks (replacement), 'tending the machine' (creating new forms of human work), and amplifying human skills (augmenting/assisting workers)—across five dimensions of meaningful work. These dimensions integrate both jobspecific (through Hackman & Oldham's, 1976 job characteristics model) and more holistic (through Lips-Wiersma & Morris', 2009 model) drivers of meaningful work. We then develop the ethical implications of our analysis by drawing on the AI4People ethical AI framework (Floridi et al., 2018) and its five principles of beneficence, non-maleficence, autonomy, justice, and explicability. We conclude with practical insights into how experiences of meaningful work will change as AI becomes more widespread and offer several directions for future research.

Al and Work: Uses and Unique Features

Current AIs constitute artificial narrow intelligence, or AIs that can undertake actions only within restricted domains, such as classifying pictures of cats (Boden, 2016). The "holy grail" of AI research is artificial general intelligence (Boden, 2016), or AIs that can perform at least as well as humans across the full range of intelligent activities. We focus only on narrow AI as it is already used across many diverse sectors, including in healthcare, judicial, educational, manufacturing, and military contexts, among many others (see Bankins & Formosa, 2021; Bekey, 2012; Walsh et al., 2019). The established use of narrow AI also allows us to draw on practical examples to ground our assessment of its effects on meaningful work. While considering the possible implications of artificial general intelligence for meaningful work is important, and we discuss this in our future research

directions, there remain persistent disagreements about when, if ever, it will be achieved (Boden, 2016). This makes it critical to examine the impacts of current AI capabilities on opportunities for meaningful work that are occurring now and in the near-term (Webster & Iyanov, 2020).

Past research demonstrates the dual positive and negative effects of technology upon aspects of meaningful work. For example, technology use can upskill workers and enhance their autonomy, but it can also deskill and serve to control them (Vallor, 2015; Mazmanian et al., 2013), with meaningfulness generally elevated in the former case (Cheney et al., 2008). Technology's positive effects can also help individuals confirm pre-existing notions of meaningful work, but its negative outcomes can require them to re-interpret and adjust those meanings as the technology's dual effects are realised, for example by providing on-demand connection to work but heightening distraction from other responsibilities (Symon & Whiting, 2019). Such dual effects remain evident in advancing forms of technology, such as workplace robotics that offer both benefits and threats to meaningful human work (see Smids et al., 2020).

These findings are critical, but their focus is on broader types of information and communication technologies, whereas we focus specifically upon AI and its implications for meaningful work. While AI use should also generate these types of dual effects, its unique features warrant specific attention. For example, compared to past technologies AI can undertake more cognitive tasks, expanding beyond 'blue collar' work in manufacturing where technology's role in replacing human labour has a long history, and into more 'white collar' forms of work (Bankins & Formosa, 2020). Further, machine learning in AIs is often driven by large amounts of data, the acquisition of which raises serious concerns about privacy, consent, and surveillance, with implications for worker autonomy (Bailey et al., 2019). Potential biases in data collection, the use of AI models built from biased data, and the resultant replication of systemic injustices (Walsh et al., 2019), as already evidenced in some AIdriven recruitment practices (Dastin, 2018), raises further concerns about the potential for one's AI-informed work to harm others. The potential for such harms is then exacerbated given the scale at which AI can be deployed. The way AIs expand opportunities to manipulate and control humans also raises important issues (Susser et al., 2019), particularly through the way it can act as an information gatekeeper for human workers (Kellogg et al., 2020). Finally, the 'blackbox' nature of the neural networks many AIs use means end-users and even AI developers cannot understand how an AI generates its outputs (Jarrahi, 2019). This can make it difficult to trust AIs, to feel competent in working alongside them, and to build responsible systems for which human workers can be held meaningfully accountable (Dahl, 2018). These features of AI have attendant consequences



for meaningful work that we will explore. We first turn to explaining the components of meaningful work and its ethical importance.

What Constitutes Meaningful Work?

Several approaches outline what constitutes meaningful work. One dominant task-based framework is Hackman and Oldham's (1976) job characteristics model (JCM), which examines how job and task design influences experiences of meaningfulness in work (Pratt & Ashforth, 2003). Other frameworks extend beyond a task focus to adopt a more "humanistic" approach (Lips-Wiersma & Morris, 2009, p. 493). For example, Lips-Wiersma and Morris (2009) suggest that meaningful work derives from finding balance between "being (true to self)-doing (making a difference)" and a focus on "self (self-actualisation)-others (serving others)". This creates the meaningful work dimensions of "developing and becoming self", "serving others", "unity with others", and "expressing one's full potential" (Lips-Wiersma & Morris, 2009, p. 501).

To adopt a holistic approach for exploring the impacts of AI on meaningful work we integrate aspects of meaningful job design from the JCM (Hackman & Oldham, 1976) with dimensions of work that facilitate the more wide-ranging enhancement of oneself through development, contribution, and connection to others from Lips-Wiersma and Morris' (2009) framework. This harmonisation generates five meaningful work dimensions that we focus our analysis upon.² The first dimension we label *task integrity*. This encompasses task identity from the JCM, or the range of tasks an individual does and the opportunity to complete a whole piece of work. This ability to undertake integrated rather than fragmented tasks then influences the extent to which workers can fully develop themselves, their capacities, and express their full potential as an integrated whole person ("developing and becoming self" and "expressing full potential" from Lips-Wiersma & Morris, 2009). The second dimension we label skill cultivation and use. This encompasses skill variety and use from the JCM, or the ability to

use and develop a range of skills at work. Like the types of tasks to which they are applied, prospects for skill utilisation then influence opportunities for growth through learning and the broader cultivation of the self and one's potential via developing, testing, and exercising a varied range of competencies ("developing and becoming self" and "expressing full potential" from Lips-Wiersma & Morris, 2009).

The third dimension is task significance (per the JCM) which connects one's work to the wider world. This dimension reflects the extent to which individuals can see how their work benefits, and contributes to the betterment of, others ("serving others" from Lips-Wiersma & Morris, 2009). The fourth dimension is autonomy (per the JCM), which reflects how freely individuals can determine their work approaches and the extent of their freedom from intrusive surveillance and monitoring. The more autonomy workers experience the greater their capacity to engage in activities like job crafting to enhance fit between individual needs and job requirements, and to undertake work that fosters self-development, moral cultivation, and that affords alignment with one's values ("developing and becoming self" and "expressing full potential" from Lips-Wiersma & Morris, 2009). The final dimension is belongingness, reflecting the ways that work can help us feel connected to a wider group to generate meaningfulness through a sense of unity with others (Bailey et al., 2019; Lips-Wiersma & Morris, 2009; Martela & Riekki, 2018). Now that we know what underpins experiences of meaning in work, we can turn to explaining the ethical dimensions of both meaningful work and AI.

The Ethics of Meaningful Work and Ethical AI

Recent philosophical discussions of meaningfulness tend to focus on what makes life itself, or the activities and relationships that compose a well-lived life, meaningful (Wolf, 2010). The paradigm of meaningless work is Sisyphus, who is condemned as punishment to repeatedly roll a rock to the top of a mountain (Camus, 1955). Sisyphus' work is boring, repetitive, simple, does not benefit others, and is not freely chosen.³ By implication, meaningful work should be engaging, varied, require the use of complex skills, benefit others, and be freely chosen. This emphasises two aspects of meaningfulness that Wolf (2010) calls subjective (do you experience work as meaningful?) and objective (is the work actually meaningful?) elements. As we take meaningful work to be "personally significant and worthwhile" (Lysova et al., 2019, p. 375), our definition is inclusive of these

³ Of course, Sisyphus' story is more complicated than this, with Camus (1955) arguing that Sisyphus finds a form of happiness in his scornful embrace of the absurdity of his condition.



¹ Pratt and Ashforth (2003) also discuss meaningfulness *at* work, or the ways leaders craft and convey organisational values to build feelings of organisational membership. To maintain a manageable scope, our analysis only examines meaning *in* work, which is largely driven by job design.

² The job characteristics model also includes feedback. We draw on the model's first three aspects as they are theorised to directly generate the psychological state of experienced meaningfulness at work, and both autonomy and belongingness are viewed in the wider literature as other critical components of meaningful work. See Parker and Grote (2022) for an assessment of technology's impact on feedback at work.

subjective (it is *personally* significant) and objective (it is worthwhile) aspects.

The Ethical Implications of Meaningful Work: Why is it Ethically Important?

Literature in business ethics and political philosophy explore the ethical significance of meaningful work (Michaelson et al., 2014). Meaningful work can be viewed as ethically significant either because it is intrinsically valuable (first basis), or because it is a constitutive element of a broader good (second basis), or because it is an instrumental good that leads to other valuable goods (third basis) (Michaelson et al., 2014). From these three bases we can see that there are good grounds for holding meaningful work to be ethically important across each of our three most used ethical theories: Kantian ethics, Virtue Theory, and Utilitarianism.

Regarding the first basis, Kantian ethical theories focus on treating people with dignity and respect as rational agents who have normative authority over their lives, and this includes imperfect duties to promote and develop the rational capacities and self-chosen ends of moral agents (Formosa, 2017). Meaningful work is ethically significant as it provides an important way to develop and exercise one's rational capacities and use them in ways that help others to meet their ends. Bowie (1998, p. 1083) identifies six features of meaningful work that explain why Kantians should care about it, including that the work is "freely entered into", "not paternalistic", "provides a wage sufficient for physical welfare", allows workers to exercise their "autonomy and independence", "develop" their "rational capacities", and promotes their "moral development". In terms of the second basis, many virtue ethicists argue that meaningful work is an integral part of flourishing as a human being. For example, Nussbaum (2011) argues that "being able to work as a human being" is a central human capability. This means being able to exercise our practical reason, use our senses, imagination and thought, have some control over our work environment, and being able to have "meaningful relations of mutual recognition with other workers" (Nussbaum, 2011, p. 34). The capability to pursue meaningful work is thus an important right and component of human flourishing. In terms of the third basis, evidence shows the positive instrumental impacts that meaningful work has on wellbeing and a range of other goods (Allan et al., 2019). This gives us good reasons to care about meaningful work for the sake of other important goods it contributes to and promotes, such as human wellbeing, that are valued on a range of ethical theories, including Utilitarianism.

Overall, according to all three of our most used moral theories there are good reasons to care about meaningful work given that it respects workers' autonomy and their ability to exercise complex skills in helping others, contributes to their wellbeing, and allows them to flourish as complex human beings. Given its ethically valuable nature, it follows that organisations have strong *pro tanto* reasons to promote, support, and offer meaningful work (Michaelson et al., 2014). Of course, *pro tanto* reasons are not indefeasible reasons, and so other considerations may outweigh them, such as improved efficiency, which means changes that lead to less meaningful work are not necessarily unethical. Further, some workers may be willing to trade off less meaningful work for other gains, such as more income or leisure time. Even so, meaningful work remains ethically important and changes that impact the amount of meaningful work for humans must be taken into ethical account, even if such considerations are not always overriding.


The Ethical Implications of AI Use

Given the ethical importance of meaningful work, more scholarship is needed to explore the potential impacts of AI upon it. The ethical significance of AI use is widely recognised and discussed (see Floridi et al., 2018; Hagendorff, 2020; Jobin et al., 2019), leading to various organisational, national, and international documents outlining ethical principles for AI deployment. However, AI's effects on meaningful work are not a focus of any of these principles. For example, Jobin et al.'s (2019) meta-analysis of ethical AI guidelines identifies 11 principles, but none mention meaningful work directly. Hagendorff's (2020) analysis also does not identify it, although related issues around the "future of employment" are discussed. An analysis by Ryan and Stahl (2020, p. 67) mentions the need to "retrain and retool" human workers who are fully replaced by AI, but this sidelines human-AI collaborations in workplaces and AI's broader impacts on meaningful work. The AI4People framework also makes no direct mention of meaningful work, but it does note the possibility of AI liberating people from the "drudgery" of some work (Floridi et al., 2018, p. 691).

While these frameworks do not mention meaningful work explicitly, we can nonetheless draw on them to identify ethical concerns that AI's impacts on meaningful work raise. To do this we draw on the AI4People ethical AI framework (Floridi et al., 2018) and its five principles of beneficence, non-maleficence, autonomy, justice, and explicability. We utilise this widely discussed framework as it emerged from a robust consensus-building program to formulate ethical AI principles. The framework's focus on the impacts of AI on "human dignity and flourishing" across its elements of "autonomous self-realisation... human agency... individual and societal capabilities... [and] societal cohesion" (Floridi et al., 2018, p. 690) also fits our focus, given that the impacts of AI on dignity, autonomous agency, social cohesion, skills and capabilities, and human flourishing all relate to our dimensions of meaningful work. The foundational principles



Fig. 1 Overview of conceptual framework



of this framework (minus explicability) have also been utilised in related work on the ethical design and deployment of broader information technologies (see Wright, 2011), which again emphasises the framework's usefulness in our context.

The five principles of the AI4People framework allow us to explore the wide-ranging impacts of AI on meaningful work. The first principle is beneficence, or the benefits AI can bring toward promoting human wellbeing and preserving human dignity in an environmentally sustainable way. Non-maleficence is about ensuring that AI does not harm humanity, and this includes not violating individuals' privacy and maintaining the safety and security of AI systems. Autonomy is about giving humans the power to decide what AI does. A linking concern between the first two principles is the use of AI, intentionally or not, to cause harm by interfering with and disrespecting human autonomy by "nudging... human behaviour in undesirable ways" (Floridi et al., 2018, p. 697). Nudging involves setting up the "choice architecture", or decision context, to intentionally attempt to push (or "nudge") people to make certain choices (Thaler & Sunstein, 2008). *Justice* is about fairly distributing the benefits and burdens from AI use and not undermining solidarity and social cohesion. Finally, explicability is about ensuring that AI operates in ways that are intelligible and accountable, so that we can understand how it works and we can require someone to be responsible for its actions. In the context of meaningful work, these principles lead us to focus on the benefits and harms that AI can bring to workers, including on their tasks, skills and social relations, the way AI might control, nudge, and manipulate workers' autonomy, the distribution of the benefits and harms AI brings, and the extent of intelligibility and accountability in AI workplace deployments.

Our overall conceptual framework is presented in Fig. 1 and our analysis is structured as follows. We first outline the impacts of AI on the five dimensions of meaningful work by analysing its effects through three pathways (outlined below) for AI deployment: replacement; 'tending the machine'; and amplifying. We then turn to the AI4People ethical framework to draw out the ethical implications of these impacts on meaningful work. This structure allows us to focus in

a systematic way on each important set of analyses, first related to AI's effects on meaningful work and then the ethical implications of this, while highlighting how these effects are often contingent on the ways in which AI is deployed.

The Effects of Artificial Intelligence on Meaningful Work

AI represents a range of technologies that can be used in many ways alongside human workers doing many different tasks. This makes it important to examine not only what tasks the AI does, but also how human workers' tasks change following AI deployment and the comparative meaningfulness of their new work. While we briefly discuss the impacts of full human replacement by AI upon meaningful work, we focus our analysis on meaningful work outcomes when humans work alongside AI.⁴ This is because such work configurations already, and will increasingly, characterise many workplaces (Jarrahi, 2018) and reflects our focus on clear current and near-term impacts of narrow AI.

Technology's Effects on Work: Three Paths

Our analytical framework adapts and expands Langlois' (2003) characterisation of how technology integrates into a work process. This structures our analysis around three key paths through which AI will shape humans' experiences of meaningful work.

In the first path, AI assumes some tasks (either simple or complex) while workers remain engaged elsewhere in the


⁴ We do not significantly detail the effects of AI fully replacing a worker because, at least currently, AI is unlikely to predominantly automate entire jobs (Chui et al., 2015). But where this does occur the impacts are clear, the unemployed worker has lost meaningful paid work until they can find another job (which may offer more opportunities for meaningful work, see Cheney et al., 2008). This also raises broader issues, beyond our scope, around other sources of meaningfulness if increasingly sophisticated AI makes paid meaningful work rarer (see Bruun & Duka, 2018).



(roughly similar) work process. This is akin to AI replacing humans in some tasks. For example, if a personalised maths learning app is introduced in a classroom, the teacher may re-focus upon other existing tasks (e.g., more time for lesson planning) or undertake new work (e.g., individualised maths coaching), but the overall work process of 'teaching' remains similar (see such examples in Acemoglu & Restrepo, 2020). We also focus on the two ends of the skills spectrum for illustrative purposes (i.e., simple and complex tasks), and acknowledge that tasks will likely involve various skills. The key difference between this path and the next is that here the replacement work undertaken by humans is not focused on managing the AI, but in the next path it is.

In the second path, AI assumes a set of tasks resulting in new human work focused on "tending the machine" (Langlois, 2003, p. 175). This is akin to creating new types of tasks for workers. We further divide 'tending the machine' into two emerging forms of work associated with managing AI: (1) what we term 'managing the machine', which generates new, complex, and interesting forms of work for humans; and (2) what Langlois (2003, p. 175) terms "minding the machine", which generates more mundane, rote, and lower-skilled work for humans. Again, we focus on two ends of a spectrum for illustrative purposes, while acknowledging that human work may exist across both categories. 'Managing the machine' reflects integrated and complex work, such as "coordination and buffering" roles (Langlois, 2003, p. 175), as well as trainer, explainer, and sustainer roles (Daugherty & Wilson, 2018). Examples include: managing the interactions between data, the wider organisation, and other stakeholders (coordination and buffering); training the AI to complete tasks and training others in AI use (training); explaining and interpreting the AI's operation and outputs to stakeholders (explaining); and ensuring the system's continued explainability, accountability, and fairness (sustaining) (Daugherty & Wilson, 2018). In contrast, 'minding the machine' work involves tasks such as "AI preparation" (sourcing, annotating, and labelling data) and "AI verification" through validating AI output (such as checking image recognition accuracy) (Tubaro et al., 2020, p. 1). This type of work tends to reflect fragmented and disconnected micro-work tasks that are often outsourced to low wage and low skill workers (Tubaro et al., 2020), leading to characterisations of "janitor work" and new digitalised forms of Taylorism (Jarrahi, 2019, p. 183).

In the third path, AI 'amplifies' or 'assists' workers by improving how human workers do their existing work



We now analyse how, through each of these three deployment pathways, AI use will impact the five dimensions of meaningful work. While individual jobs could experience elements of all three paths (e.g., some replacing, some 'tending the machine' work, and some amplifying) and some overlap may occur (e.g., AI replacing a rote human task also assists the worker), we discuss each path as distinct for analytical purposes. The ethical implications of these impacts are then assessed in the subsequent section.

Task Integrity and Skill Cultivation and Use

Workers' tasks can range from being highly fragmented to being highly integrated, and the diversity of skills they can activate will also vary as a result. Both of these aspects generate opportunities to achieve and develop one's abilities and potential through work (Lips-Wiersma & Morris, 2009). As the nature of what a worker does (i.e., tasks) strongly impacts what they need to do that work (i.e., skills), we discuss these two dimensions together.

First, we consider the path of AI taking over some tasks while leaving workers engaged in other work. The tasks the AI assumes could be *simple* or *complex* (or anything in between), but the predominance of narrow AI means it is mainly deployed to replace humans in specific narrow tasks. An espoused benefit of AI is its ability to undertake simple tasks that are often boring and unchallenging for humans, such as collating information for meetings (Pulse + IT, 2020) or assessing fruit quality (Roberts, 2020). Deploying AI in this way is unlikely to generate significant feelings of marginalisation from a wider work process due to the simple nature of the tasks it is assuming, particularly when the human takes on other comparable or more interesting work. This should result in neutral or improved perceptions



⁵ We acknowledge that other forms of new human work are also likely to emerge (see Acemoglu & Restrepo, 2020), but its nature remains speculative. New work associated with AI management already exists or is emerging, aligning with our focus on near-term work implications of AI.

of task integrity and may free workers' time to engage in more learning and development.

However, when AI assumes more complex and significant tasks then its implications, both positive and negative, may be more profound. For example, in human resource management an AI can shortlist candidates to progress to interviews based on natural language processing of applications (Bankins, 2021; Leicht-Deobald et al., 2019). Shortlisting applicants can be a complex and significant component of the recruitment and selection process. Using AI for this task could then degrade workers' experiences of task integrity as they no longer undertake a significant part of a work process, assuming this work is not comparably replaced. Shifting workers to other more rote tasks, despite adding work that maintains their level of involvement in the work process, is also likely to compound feelings of reduced task integrity, as the worker moves from undertaking more significant to less significant work. This can also limit the scope for workers to develop and express their full capabilities at work and reduce their opportunities for growth.

In contrast, if workers shift to new but similarly complex or even more significant tasks elsewhere in the work process, then this should support task integrity as the worker continues to contribute meaningfully to work outcomes. For example, the AI 'AlphaFold' developed by DeepMind is designed to automate and accelerate the process of determining protein structures, an important step in developing new treatments for human diseases (Hassabis & Revell, 2021). While AlphaFold can assume significant tasks previously done by human scientists (i.e., determining protein structures) this should positively impact, or at least have a neutral effect, on task integrity if it allows scientists to re-focus their work efforts on other important aspects of their broader goal of curing diseases. However, there remain risks to AI being used in this way. Continuing with this example, if scientists have trained for many years to do the experimental work that AlphaFold can now do more quickly and accurately, this generates significant risks for their ability to exercise their full capacities, demonstrate their mastery, and utilise the skills they have invested years in developing to reach their full potential.

Changes in skill cultivation and use due to technology replacing either simple or complex tasks also raises deskilling concerns, whereby skilled human work is offloaded to machines resulting in skill loss (Vallor, 2015). Ethically, it is critical to establish whether the human skills lost (i.e., offloaded to machines) are important and whether they can be exercised and maintained through other forms of work or in other life domains (Michaelson et al., 2014; Wolf, 2010). As simple and rote work generally requires basic skills that can be cultivated elsewhere or are not significant, there is limited scope for significant deskilling in this case. However, complex tasks generally require complex skills, such

as judgement, intuition, context awareness, and ethical thinking. From a deskilling perspective, these types of skills are particularly ethically problematic for workers to risk losing. This means when workers are left with fewer overall complex and significant tasks following AI deployment, then their ability to cultivate and use important skills will likely decrease, negatively impacting this dimension of meaningful work.

It is worth noting that where replacement involves AI assuming a worker's whole job, for example where the job is constituted entirely of simple and rote tasks that are most susceptible to full automation (Gibbs, 2017), this will likely lead to unemployment (if redeployment is not possible). This effectively removes, at least temporarily, paid meaningful work from that worker's life and poses the greatest risk to the ability to experience meaningful work. This also provides the conditions for a wide range of skills to be lost or degraded, as well as having significant negative impacts on important self-attitudes, such as feelings of self-respect and self-worth (see Selenko et al., 2022 for work on AI use and employees' sense of identity). This case also raises broader political questions about how society should deal with such a scenario should it become more widespread (Hughes, 2014). While these questions are beyond our focus here, we do highlight them in our discussion of future research directions.

Second, we consider the path of workers 'tending the machine', whether in 'managing' or 'minding' forms. 'Managing the machine' work should enhance what Bourmault and Anteby (2020, p. 1453) term "administrative responsibility", through offering a wider scope and variety of duties. This should enhance task integrity where the shift to coordination and buffering work provides opportunities for integrated and challenging activities across training, explaining, and sustaining roles through supervisory work, technology oversight, exceptions management, and cross-functional coordination of entire work processes. Such coordination and buffering work will also require the development of flexible and wide-ranging skill sets (Langlois, 2003), supporting skill cultivation and use and more broadly widening and deepening one's ability to learn, achieve, and develop at work.

In contrast, rather than generating more complex and interesting human work, 'minding the machine' produces a "more benignant role for humans" through more mundane and rote tasks (Langlois, 2003, p. 174). This would reduce task integrity as workers become more distanced from their work outcomes. The generally repetitive and fragmented nature of 'minding the machine' work also suggests its associated skills are low and narrow, offering little opportunity for varied skill cultivation. Such AI "janitor work" (Jarrahi, 2019, p. 183) risks degrading workers' abilities to meaningfully develop their capabilities and reach and express their



full potential at work, leading to lower levels of meaningfulness on this dimension.

Third, when AI *amplifies* workers' abilities to do their current tasks, positive impacts on task integrity and skill cultivation and use should ensue. For example, in the policing domain, machine learning technologies can collate previously disparate data sources to analyse characteristics and histories of domestic violence victims and perpetrators to better predict, compared to current human-driven systems, repeat attacks and better prioritise preventative actions (Grogger et al., 2020).⁶ In such cases, experiences of task integrity are likely to remain consistent or improve as AI supports workers to better complete their tasks and achieve work goals. Skill cultivation and use should remain neutral or improve as it is likely that workers, while maintaining their current skills, will need to develop new ones to interpret and integrate AI output into their decision making.

However, a feature of AI that may constrain skill use across all three paths is its 'blackbox' nature (Boden, 2016). While AI designers are developing ways to improve lay person interfaces, the use of 'blackbox' (or unexplainable) AI in workplaces may degrade workers' skill cultivation, use, and feelings of competence. For example, where workers are highly reliant on the decision making of an AI, they may feel lower levels of competence in their use of it due to little understanding of its functioning. This effect will likely be more acutely felt where workers are expected to understand and explain what the AI is doing. Poor explainability can also create opaque chains of accountability for decisions informed by AI (Dahl, 2018) and this risks making workers overly dependent on an AI that they cannot comprehend.

Task Significance

Task significance means employees see their work as having positive impacts (Grant, 2008) through their service to others (Lips-Wiersma & Morris, 2009), within or outside the organisation (Hackman & Oldham, 1975). Task significance is influenced by how employees assess their *job impact on* and *contact with* beneficiaries (Grant, 2007). Job impacts on beneficiaries are shaped by the dimensions of magnitude, scope, frequency, and focus (i.e., preventing harms or promoting benefits) (Grant, 2007). Contact with beneficiaries is shaped by the dimensions of frequency, duration, physical proximity (including virtual proximity), depth, and breadth of contact (Grant, 2007). Given the range of these dimensions, workers' assessments of task significance can

⁶ Although in practice such predictive policing systems have been shown to risk biased outcomes against minority groups, driven by over-representation of those groups in policing statistics (Berk, 2021). We discuss these issues in a later section.



be complex. Evidence suggests that employees can derive task significance from even objectively rote, mundane, and low skill work (the objective element of meaningful work), when that work is framed in the right way (the subjective element of meaningful work). For example, Carton (2018, p. 323) shows that when leaders at NASA carefully framed the space agency's goals, workers could connect work such as "mopping the floor" to "helping put a man on the moon". However, carrying out impactful tasks without opportunities for "personal, emotional connections to the beneficiaries of those tasks" can impede overall experiences of task significance (Grant, 2007, p. 398; Bourmault & Anteby, 2020). This means that both job impact on beneficiaries and contact with them are important to assess. Given this, and following Grant (2007), we suggest that employees' global assessments of the impact of AI on their jobs, rather than on specific tasks, is most relevant when assessing perceptions of task significance.

First, we consider the path of AI taking over *simple* or complex tasks while workers remain engaged elsewhere. Given our focus here at the job level, if only some simple tasks are assumed by an AI this should have limited impact on task significance, assuming the remaining or new tasks provide opportunities for workers to positively impact and connect with beneficiaries. In contrast, when AI assumes more *complex tasks*, these are likely significant to an individual's overall assessments of task significance. This may lead to more extensive and complex sensemaking of this change. To see this, we draw on construal-level theory (CLT), which describes the way individuals cognitively represent people or events at either higher or lower levels of abstraction (Trope & Liberman, 2003). Higher levels of abstraction involve "mental representations that are relatively broad, inclusive, (and) general", such as higherlevel goals or principles (Wiesenfeld et al., 2017, p. 368). Lower levels of abstraction involve "applying relatively specific, detailed, and contextualised representations", such as focusing on lower-level actions to achieve higher-level goals (Wiesenfeld et al., 2017, p. 368). Returning to the earlier AlphaFold example, at a higher level of construal workers may perceive improved task significance regarding job impact as the AI is significantly contributing to the higher-level goal of treating diseases. This could facilitate higher perceptions of magnitude, scope, and frequency of positive impact. At a lower level of construal, the worker may then ask: "but what am I doing to help meet this goal?". If workers can re-focus on other comparatively significant tasks in the work process, they should experience higher task significance as the AI helps advance the field toward reaching the overarching goal and the worker continues to meaningfully contribute toward that goal. However, where the remaining or new tasks fail, at lower levels of construal, to deliver at least the same experiences of task significance as before, then one's perceived ability to 'serve others' is likely to degrade overall.

In cases of both simple and complex task change, where AI is used in ways that have sub-optimal, biased, unjust, or harmful outcomes for end users, this could also decrease workers' perceptions of task significance. For example, where AI provides facial recognition and predictive policing data for law enforcement agencies and the AI's outputs are biased against minority groups, then workers may see reduced task significance given their organisations' connections to negative outcomes (via the negative magnitude, scope, and frequency dimensions of job impact). Implicating workers in injustices and harms perpetrated by an AI, through their involvement with or responsibility for the technology, can particularly diminish the experience of serving others and the autonomous ability to act in alignment with one's values and morals (related to 'developing and becoming self'), degrading overall work meaningfulness.

Second, we consider the 'tending the machine' path. In terms of 'managing the machine', the heightened administrative responsibility associated with such work (Bourmault & Anteby, 2020) should improve task significance through job impact and more opportunities to benefit others, given the expansive duties this work entails (i.e., enhanced scope of impact). However, such work can distance workers from those they serve, diminishing feelings of direct "personal responsibility" (Bourmault & Anteby, 2020, p. 1453) or feelings of having a direct and significant impact on the lives of others, which can reduce task significance. For example, when replaced by autonomously driven trains and moved to 'managing the machine' work, metro train drivers experienced enhanced administrative responsibility but diminished personal responsibility, alongside lower task significance overall, as they were no longer directly responsible for commuters' safety (Bourmault & Anteby, 2020).

In terms of 'minding the machine' work we suggest that task significance will generally be reduced. This is because such fragmented work means workers may have little idea of the point of their labour and its impacts, potentially limiting all job impact dimensions. As they may also be working in isolation from others because of outsourcing (Tubaro et al., 2020), potentially limiting all contact with beneficiaries, this further disconnects workers' tasks from the end user benefits generated, eroding task significance.

Third, when AI *amplifies* a worker's abilities this should have significant and positive implications for task significance, particularly through the magnitude and focus dimensions of job impact and the duration and depth dimensions of contact with beneficiaries. Here, the AI is not focused on substantially changing the range of tasks in a work process, but rather on improving something that humans were already doing in that process, leading to better outcomes for beneficiaries. Drawing on earlier amplification examples,

where AI can support police officers by collating and analysing new data sources to help them better prevent incidences of domestic violence (assuming it does not do so in unfair or biased ways), then this should heighten perceptions of both being able to achieve higher-level goals (e.g., preventing crime) and seeing the importance and connection of lower-level tasks to reaching that goal (e.g., through interpreting better predictive analytics). Use of AI in this way can also help reduce human biases in decision making, such as through building fairness principles into AI systems (see Selbst et al., 2019). In recruitment, for example, AI can limit the impact of unconscious human biases and various other human constraints on rational decision making by assessing all candidates' applications against standard criteria and providing auditable, transparent, and explainable decision trails (see Hagras, 2018; Bankins et al., 2022). This path demonstrates that a significant potential benefit of AI is that it can elevate humans' abilities to address complex problems, enhance the impact of their work, and thus better serve others, through its analysis of large datasets to identify novel insights.

Autonomy

Autonomy means self-rule. Individually, that means being able to do what you really want to do. In addition to the freedom from interference needed to rule yourself, autonomy is also commonly taken to include competency (i.e., you have the skills and capacities needed to rule yourself) and authenticity (i.e., your ends are authentically your own and not the result of oppression, manipulation, or coercion) conditions (Formosa, 2021). In the workplace, autonomy refers to "the degree to which the job provides substantial freedom, independence, and discretion to the individual in scheduling the work and in determining the procedures to be used in carrying it out" (Hackman & Oldham, 1976, p. 258). AI's impact on individuals' autonomy is a key issue for the ethical AI literature. A particular concern is that ceding authority to AI diminishes human autonomy (Floridi et al., 2018). However, potential benefits for human autonomy can also accrue from increasing AI's autonomy. We assess these different impacts of AI at work as either promoting or diminishing autonomy across competency and authenticity conditions.

In terms of promoting human autonomy, this depends on what work the AI assumes but also, and more importantly, on what work takes its place and what control and input workers have over AI deployment. When AI assumes *simple* or *complex tasks* that workers find boring or repetitive, then this potentially promotes autonomy by freeing up time for workers to build their autonomy competencies through doing other more challenging or authentic work. For example, if an AI prioritises a worker's emails so that she only sees those requiring a response, this may free her to work on other



more valuable tasks. In terms of 'managing the machine', this path could promote autonomy if new work is more skilful and engaging than the work it replaces, and if workers have a degree of control over how that work is done. Where AI amplifies workers by giving them more power and useful information, then this can improve worker autonomy by helping them to better achieve their self-given ends.

In terms of diminishing human autonomy, these impacts are partly the converse of the above. In our first path, if complex, interesting, and creative tasks that workers want to do are assumed by AIs, this potentially diminishes autonomy. There may also be good reasons why humans should remain engaged in certain complex tasks and decisions, such as due to their moral complexity. This means that where AI assumes these tasks it can diminish human achievement of valuable ends, degrade important human skills, and limit opportunities for moral development (Lips-Wiersma & Morris, 2009). For example, when we delegate to AI decisions regarding ethically sensitive aspects of human resource management, the skills associated with that work can degrade and thereby diminish important autonomy competencies. AI can also make our autonomy more vulnerable by making us dependent on it, which means our autonomy can diminish if access to the technology is removed. Across the 'tending the machine' path, through 'managing the machine' work AI can diminish worker autonomy by filtering and potentially restricting the information that is made available for humans to view and use (Kellogg et al., 2020). Such constraints can limit the ability for workers to authentically develop themselves and their capabilities at work. Broader autonomy concerns also exist with 'minding the machine' work, which is itself mundane and boring, making workers feel like a 'slave to the machine' (Engel, 2019) and thereby experiencing diminished autonomy at work.

Across all paths, a more pernicious threat to autonomy may exist through surveillance and manipulation by AI. This reflects what Foucault calls the rise of a "surveillance society", which seeks to control bodies through making people feel permanently monitored (Abrams, 2004). When people are surveilled they tend to feel constrained and act in less authentic and autonomous ways (Molitorisz, 2020). The use of AI to surveil workers will likely have similar impacts and can be a way for employers to use their power to exert control over employees. For example, the use of AI-powered cameras to surveil Amazon delivery drivers could make them more self-conscious in their trucks, which could lead them to feel more constrained and unable to act autonomously (Asher-Schapiro, 2021). A similar example is when AI is implemented to monitor online meetings and measure whether workers are engaged and contributing to the discussion (see Pardes, 2020), which could lead to stress and inauthentic behaviour. Such monitoring could also result in workers engaging in intentional "deviance" to challenge the control of surveillance (Abrams, 2004), by trying to "game" the AI by matching or openly flouting what the AI is expecting in terms of eye contact and body language (Pardes, 2020), or finding other ways to operate outside the gaze of the surveillance system.

Belongingness

Belongingness refers to "the meaningfulness of working together with other human beings" (Lips-Wiersma & Wright, 2012, p. 673). Across all our paths, we argue that AI may impact workers' belongingness in two main ways: through generating the conditions for more or less meaningful connections and a sense of unity with others; and through its implementation creating differences across workers that undermines solidarity.

In terms of the first way, where AI assumes tasks that may otherwise have required in-person and face-to-face interaction with other workers or customers, this can create less human contact in the workplace. For example, where an AI chatbot allows workers to access information previously provided by a human worker, this lessens that worker's interactions with other humans and reduces opportunities for forming connections with others that are the bedrock for generating a sense of belonging (Seppala et al., 2013). In contrast, AI use may increase opportunities for human interaction, for example through 'managing the machine' work where workers are responsible for supervising AI deployment that requires extensive human-to-human training.

In terms of the second way, a key concern in the ethical AI literature is how AI use may disproportionately and negatively affect lower-skilled and lower-paid workers, while its benefits may disproportionately accrue to those with higher skills and wages (Ernst et al., 2018), effectively creating new types of workplace in-groups and out-groups. For example, many of the negative impacts of AI at work, such as surveillance and simplistic 'minding the machine' work, will tend to fall on less skilled 'blue collar' workers, whereas more of the amplifying and autonomy-enhancing benefits associated with taking on even more interesting and engaging work will tend to fall to already privileged workers. This creates justice concerns around how the benefits and burdens of AI in workplaces are being distributed, potentially undermining solidarity between those who benefit from AI's introduction and those who do not. For example, in a call centre context an AI may be used to monitor and evaluate the calls of every call centre operator. Such heightened surveillance may be perceived by operators as intrusive and diminishing their autonomy. However, using AI in this way may amplify the work of quality assurance staff in the same organisation, providing them with more information and assisting them in better training and managing operators. This shows how AI may generate distinct groups experiencing very different



impacts, such as being viewed as unnecessary surveillance by some but as an amplifying source of information by others. Such outcomes particularly threaten the ability to create a sense of belongingness and shared values (Lips-Wiersma & Morris, 2009), which underpins the 'unity with others' dimension of meaningful work.

Ethical Implications: Al and Meaningful Work

We have analysed how the three paths of AI deployment may enhance or diminish opportunities for meaningful work across five dimensions. We now surface the ethical implications of this analysis via the five principles of the AI4People ethical AI framework (Floridi et al, 2018): beneficence; non-maleficence; autonomy; justice; and explicability. As with any principlist framework there are potential conflicts and tensions between principles (Formosa et al., 2021). For example, there may be benefits for some from AI deployment (beneficence) while others suffer harm (non-maleficence) or interference with their autonomy. As identified earlier, the provision of meaningful work is not always the only or most important ethical value at stake, and so less meaningful work may not be ethically worse overall if there are other ethical benefits, such as improved wellbeing for others through higher productivity.

To assess ethical implications, we synthesise and summarise how the three paths (replacing, 'tending the machine', and amplifying) support or limit experiences of meaningful work and so contribute to, or diminish, meeting the AI4People principles. We summarise these impacts in Table 1 across the five ethical principles (beneficence and non-maleficence are combined in the Table as the latter reflects the converse of the former), while noting the main deployment pathways through which these impacts occur.

In terms of the beneficence principle, there can be significant benefits for employees when AI use supports the various dimensions of meaningful work. When AI amplifies a worker's skills it can support them to complete their tasks, undertake more complex tasks, and utilise higher-order thinking and analysis skills (task integrity and skill cultivation and use). It can also afford workers the opportunity to achieve better outcomes and enhance the positive impact of their work on beneficiaries (task significance), give them more control over their work through improved access to information (autonomy), and potentially generate new connections with other workers and stakeholders (belongingness). Similarly, when AI assumes some *simple* or *complex* tasks and the human worker can re-focus on other important and challenging tasks in the work process, then positive experiences across all dimensions of meaningful work should be maintained or improved. 'Managing the machine'

work can also improve meaningfulness through a wider scope of enriched work (task integrity and skill cultivation and use) and a wider positive job impact within and outside the organisation (task significance), as well as greater interaction with a range of stakeholders through coordination and supervisory work (belongingness).

In terms of the non-maleficence principle, we also show the harms that AI can create when it is deployed in ways that lead to less (or no) meaningful work, or other related harms. Two paths generate greatest risk of harms through significantly reducing experiences of meaningful work. First, when AI replaces some tasks, the risk of degraded task integrity, deskilling, reduced task significance, and constrained autonomy is greatest when it assumes more *complex* tasks and the worker is not afforded any new comparable or more interesting work. This is because complex tasks generally constitute a large and significant part of the work process and undertaking them exercises a range of important skills. Being removed from such work can also distance workers from the output of their labour and lower perceptions of beneficiary impact. In the worst case, it could involve the complete loss of paid meaningful work where AI replaces whole jobs, which removes workers from important social relationships and denies them the opportunity to skilfully utilise their talents to help others. Second, 'minding the machine' work, as we have characterised its fragmented, piecemeal, and micro-work nature, threatens these same aspects of meaningful work and feelings of belongingness when work is outsourced to disconnected workers. Other paths can also generate harms, but arguably at lower levels. For example, we identified that while 'managing the machine' work may increase meaningful work experiences overall through heightened administrative responsibility, it can lessen feelings of task significance by increasing distance between workers and their beneficiaries and reducing feelings of personal responsibility.

In terms of the autonomy principle, across each path we show how autonomy is supported when AI is used to free up humans to focus their time on other more valued tasks, allows them to develop new or enhanced autonomy competencies, and gives them more control over their work. In particular, the task replacement, 'managing the machine', and amplifying paths that afford employees access to better data and information, the opportunity to engage in more interesting work, and exercise more control over how their work is done, can all promote autonomy as a dimension of meaningful work. However, many of these positive impacts also depend on whether workers have input into how AI is deployed in their organisations. A particular risk to autonomy is the use of AI to surveil and monitor, which can undermine authenticity and encourage workers to align their behaviours with the AI's implicit expectations or seek ways to subvert or avoid its control.



Table 1 Ethical impacts of AI for enhancing or diminishing meaningful work Ethical aspects of meaningful work (AI4 How AI could enhance meaningful work (and How AI could diminish meaningful work (and People principles) main pathways) main pathways) Beneficence & Non-maleficence Main pathways: Replacing, amplifying, & Main pathways: Replacing & minding the managing the machine machine • Loss of work for humans • Less boring and repetitive human work • Increased opportunities for new and/or more • Reduced human role in the work process challenging human work Human deskilling • Enhanced human learning, skills, and · More boring and fragmented 'minding the development machine' work • Augmenting and assisting human workers · Less belonging and less human interaction • Facilitating higher positive impacts on Autonomy Main pathways: Replacing, amplifying, & Main pathways: All pathways, especially mindmanaging the machine ing the machine • Improved autonomy competencies (e.g., • Reduced autonomy competencies (e.g., through more time for valuable work) through deskilling) • More control and power for workers through • Surveillance greater access to information • Manipulation & nudging (e.g., controlling human behaviour) • Employers exerting more power and control over workers · Worker vulnerability & dependence on AI · Inauthentic behaviours (e.g., acting more selfconsciously) • Resistance to AI and greater deviance (e.g., 'gaming' the AI) Main pathways: All pathways Justice Main pathways: Replacing & amplifying • More data-driven choices • Unfair distribution of benefits and burdens • Less human bias in decision making (e.g., lower-skilled workers suffer more bur-· Fairer decision making dens and receive less benefits) Undermining solidarity · Implicating workers in injustices resulting from AI bias Explicability (explainability & accountability) Main pathways: Managing the machine & Main pathways: Replacing, amplifying, & amplifying managing the machine . Informating (i.e., improved access to infor-· Feelings of incompetence due to poor AI mation) explainability • Greater transparency in decision making • Unclear chains of accountability when AI is • Upskilling in understanding AI involved in decision making

The justice principle centres on ensuring fair, just, and non-discriminatory outcomes from AI and requires a focus on how the benefits and burdens of AI use are distributed. For example, the amplifying path generally achieves strongly positive outcomes for meaningful work, but there is evidence that such benefits are disproportionately allocated to already privileged workforces (i.e., higher-skilled and higher-paid workers). In contrast, the 'minding the machine' path generally achieves strongly negative outcomes for meaningful work, but such burdens tend to disproportionately impact less privileged workforces (i.e., lower-paid and lower-skilled workers). Lower-skilled workers are also more likely to have their entire jobs replaced by AI (Gibbs, 2017). This uneven distribution raises important justice concerns and can undermine solidarity and feelings of belongingness within and across work groups. However, AI can also be deployed to promote justice, which can positively impact task significance. For example, when AI is used to minimise bias and maximise evidence-based decision making through giving workers access to new data-driven insights (such as through amplification, 'managing the machine', or replacing complex tasks paths), this promotes fair outcomes while also enhancing task significance through a greater positive impact on beneficiaries. But the converse also holds when the justice principle is threatened by an AI trained on biased datasets and deployed in workplaces where it generates unjust outcomes that can decrease task significance and implicate workers in injustices.

Finally, the *explicability principle* relates to the explainability, transparency, and accountability of AI. In paths where AI plays a significant role alongside human workers, such as the *amplifying*, '*managing the machine*', and the replacement of *complex tasks* paths, an inability of workers to understand an AI's operation, particularly where they



are highly reliant upon it and are accountable for it, can constrain skill use and feelings of competence. This could potentially undermine the benefits AI may otherwise bring. This suggests that training workers not only in what AI does but also how it does it, and making chains of accountability clear, will be important for supporting experiences of meaningfulness at work.

Practical Implications

Organisational use of AI can reap many benefits through improved service range and quality, efficiency, and profitability. However, the ethical deployment of AI requires weighing up its many costs and benefits. We help articulate some of those costs and benefits for workers in terms of AI's impacts on meaningful work. Practically, this is important because some authors suggest an emerging trend is for organisations to use AI for full automation (Acemoglu & Restrepo, 2020), without also considering opportunities to use it for enhancing human work, and then poorly preparing their workforces for the changes that AI use entails (Halloran & Andrews, 2018). For organisations we highlight those pathways, such as 'minding the machine' work, that are likely to significantly limit opportunities for meaningful work, which implies that other considerations such as efficiency benefits must strongly outweigh the harms to workers that AI used in this way can generate, in order to justify its use. We also highlight that, when considering meaningfulness, it is insufficient to focus only on the AI itself, as the implications of its deployment are strongly driven by what work remains for humans, which is something that organisations can directly influence and decide. Overall, we offer guidance on how organisations can maintain or build opportunities for meaningful work when they implement AI and point leaders toward specific areas for intervention to support meaningful work experiences. For example, task significance is critical for meaningful work (Grant, 2007), yet the ways AI can distance workers from beneficiaries threatens these experiences. However, there are ways in which organisations can remedy this, such as by sharing end users' positive stories with workers (Grant, 2008).

Future Research Directions

Although we did not frame explicit propositions from our conceptual work, there are several relationships we suggest warrant empirical examination. For example, assessing whether AI performing *simple tasks* enhances task integrity, but AI performing *complex tasks* degrades this

dimension and perceptions of task significance, and examining whether the 'managing the machine' and amplifying paths enhance task integrity and skill cultivation and use overall, but 'minding the machine' work diminishes these aspects. We also suggest several contingencies will affect these relationships, such as what other or new work employees do following AI implementation (task-related factors) and how aspects of the technology, such as its explainability or potential for bias, shape workers' experiences of it (technology-related factors).

While we adapted Langlois' (2003) work to develop our three pathways, these may manifest in different ways and will likely overlap. Future research could explore how each path operates in workplaces, how they may differ from our conceptualisation, and whether there are other path configurations to AI deployment that our framework does not capture. There may also be nuances within pathways that warrant investigation. For example, Jarrahi (2018, p. 3) suggests that advances in AI could create new forms of "human–machine symbiosis" that result in "both parties (becoming) smarter over time". This could generate new forms of human skills, tasks, and perhaps whole jobs that have not yet been imagined, with implications for meaningful work.

Another area for future work is examining how leaders construct and influence subjective perceptions of the meaningfulness of work, particularly through the values, strategies, and vision that underpin how they implement AI (Pratt & Ashforth, 2003). For example, if an organisation is focused on full automation and replacing human workers, it will likely deploy AI toward this end and degrade opportunities for meaningful work. But if leaders adopt multistakeholder governance approaches that support ethical AI deployment (Wright & Schultz, 2018), such participatory practices may enhance perceptions of meaningful work following AI deployment.

Finally, while we centred our analysis on the meaningful work implications of narrow AI, future work could utilise conceptual tools such as thought experiments (Bankins & Formosa, 2020) and work on posthumanism (Gladden, 2016) to prospectively analyse the impacts of potential future forms and deployments of more advanced AI. For example, developments in virtual and augmented reality are creating movements toward a metaverse, or a persistent form of virtual world that is accessible through various devices and that people combine with their existence in the physical world (Ravenscraft, 2021). Such technologies have the potential to transform the nature of social interactions and thus impact the belongingness dimension of meaningful work. Likewise, advances in natural language processing and speech interfaces could result in workers having multiple "digital assistants" (Zhou et al., 2021, p. 258), which will




impact the nature of workers' tasks and relationships and the skills they will require in the future.

Extending even further, artificial general intelligence (AGI) would constitute "a new general-purpose technology" (Naudé & Dimitri, 2020) that has been predicted to pose existential threats such as eradicating large swathes of human work (Bruun & Duka, 2018) and even risking humanity's annihilation (Torres, 2019). The reality of such technologies would inevitably lead to more extreme conclusions for the future of meaningful work than we have generated here through our focus on narrow AI, as they would likely render all but our replacing path largely obsolete. The possibilities of such technologies may therefore lead us back to substantive discussions on the value of work generally, and what forms of human work we believe must be preserved or newly created no matter what technologies are developed. This also raises questions of what broader social changes, such as increased volunteering, provision of other forms of meaningful activity, or a Universal Basic Income (Hughes, 2014), will be required to cushion negative impacts should AGI deployment ever become a reality. It also augurs the potential for heavier regulation of the development and use of AI (and potentially AGI) to maintain meaningful forms of human employment, and to place limits on where, how, and why AI is used. However, at this point the discussion relies on largely technical questions about whether AGI is indeed possible (Boden, 2016). In the meantime, the impacts of narrow AI on meaningful work are ones we need to address here and now.

Conclusion

This paper focused on a neglected aspect of the ethical implications of AI deployment, namely the impacts of AI on meaningful work. This is an important contribution as the ethical AI literature, while focused on the impacts of unemployment resulting from AI, needs to also attend to the impacts of AI on meaningful work for the remaining workforce. Given the ethical importance of meaningful work and its considerable impacts on human wellbeing, autonomy, and flourishing, this is a significant omission that we help to remedy. We have done so by examining the impacts of three paths of AI deployment (replacing tasks, 'tending the machine', and amplifying) across five dimensions of meaningful work (task integrity, skill cultivation and use, task significance, autonomy, and belongingness). Using this approach, we identify specific ways in which AI can both promote and diminish experiences of meaningful work across these dimensions and draw out the ethical implications of this by utilising five key ethical AI principles. Finally, we offer practical guidance for organisations by articulating the ways that AI can be implemented to support



Acknowledgements The authors would like to sincerely thank the Special Issue Guest Editors, their Action Editor Associate Professor Luke Fletcher, and the anonymous reviewers for their insightful and constructive feedback during the review process.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Declarations

Conflict of interest The authors have no conflicts of interest to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Abrams, J. J. (2004). Pragmatism, artificial intelligence, and posthuman bioethics: Shusterman, Rorty, Foucault. *Human Studies*, 27(3), 241–258.

Acemoglu, D., & Restrepo, P. (2020). The wrong kind of AI? Artificial intelligence and the future of labour demand. *Cambridge Journal* of Regions, Economy and Society, 13, 25–35.

Allan, B. A., Batz-Barbarich, C., Sterling, H. M., & Tay, L. (2019). Outcomes of meaningful work: A meta-analysis. *Journal of Management Studies*, 56(3), 500–528.

Asher-Schapiro, A. (2021). Amazon AI van cameras spark surveillance concerns. News.Trust.Org. https://news.trust.org/item/ 20210205132207-c0mz7/

Bailey, C., Yeoman, R., Madden, A., Thompson, M., & Kerridge, G. (2019). A review of the empirical literature on meaningful work: Progress and research agenda. *Human Resource Development Review*, 18(1), 83–113.

Bankins, S. (2021). The ethical use of artificial intelligence in human resource management: A decision-making framework. *Ethics and Information Technology*, 23, 841–854.

Bankins, S., & Formosa, P. (2020). When AI meets PC: Exploring the implications of workplace social robots and a human-robot



- psychological contract. European Journal of Work and Organizational Psychology, 29(2), 215–229.
- Bankins, S., & Formosa, P. (2021). Ethical AI at work: The social contract for artificial intelligence and its implications for the workplace psychological contract. In: M. Coetzee & A. Deas (Eds.), Redefining the Psychological Contract in the Digital Era: Issues for Research and Practice (pp. 55–72). Springer: Switzerland.
- Bankins, S., Formosa, P., Griep, Y., & Richards, D. (2022). AI decision making with dignity? Contrasting workers' justice perceptions of human and AI decision making in a human resource management context. *Information Systems Frontiers*, 24(3), 857–875.
- Bekey, G. A. (2012). Current trends in robotics. In P. Lin, K. Abney, & G. A. Bekey (Eds.), *Robot ethics* (pp. 17–34). MIT Press: Cambridge, Mass.
- Berk, R. A. (2021). Artificial intelligence, predictive policing, and risk assessment for law enforcement. Annual Review of Criminology, 4(1), 209–237.
- Boden, M. A. (2016). AI. Oxford University Press: UK.
- Bourmault, N., & Anteby, M. (2020). Unpacking the managerial blues: How expectations formed in the past carry into new jobs. *Organization Science*, *31*(6), 1452–1474.
- Bowie, N. E. (1998). A Kantian theory of meaningful work. *Journal of Business Ethics*, 17, 1083–1092.
- Bruun, E., & Duka, A. (2018). Artificial intelligence, jobs and the future of work. *Basic Income Studies*, 13(2), 1–15.
- Camus, A. (1955). The myth of Sisyphus and other essays. Hamish Hamilton.
- Carton, A. M. (2018). I'm not mopping the floors, I'm putting a man on the moon: How NASA leaders enhanced the meaningfulness of work by changing the meaning of work. *Administrative Science Quarterly*, 63(2), 323–369.
- Cheney, G., Zorn Jr, T. E., Planalp, S., & Lair, D. J. (2008). Meaningful work and personal/social well-being organizational communication engages the meanings of work. *Annals of the International Communication Association*, 32(1), 137–185.
- Chui, M., Manyika, J., & Miremadi, M. (2015). *The four fundamentals of workplace automation*. McKinsey. http://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/four-fundamentals-of-workplace-automation
- Dahl, E. S. (2018). Appraising black-boxed technology: The positive prospects. *Philosophy & Technology*, 31, 571–591.
- Dastin, J. (2018, October 11). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. https://www.reuters. com/article/us-amazon-com-jobs-automation-insight-idUSK CN1MK08G
- Daugherty, P. R., & Wilson, H. J. (2018). *Human + Machine: Reimagining work in the age of AI*. Harvard Business Review Press.
- Engel, S. (2019). Minding machines: A note on alienation. *Fast Capitalism*, 16(2), 129–139.
- Ernst, E., Merola, R., & Samaan, D. (2018). The economics of artificial intelligence. International Labour Organization. https://www.ilo.org/wcmsp5/groups/public/dgreports/cabinet/documents/publication/wcms_647306.pdf
- Floridi, L., et al. (2018). AI4People An ethical framework for a good AI society. Minds and Machines, 28(4), 689–707.
- Formosa, P. (2017). Kantian ethics, dignity and perfection. Cambridge University Press: Cambridge.
- Formosa, P. (2021). Robot autonomy vs human autonomy: Social robots, artificial intelligence (AI), and the nature of autonomy. *Minds and Machines*, *31*, 595–616.
- Formosa, P., & Ryan, M. (2021). Making moral machines: Why we need artificial moral agents. *AI & Society*, *36*, 839–851.
- Formosa, P., Wilson, M., & Richards, D. (2021). A principlist framework for cybersecurity ethics. *Computers & Security*, 109, 102382.

- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.
- Gibbs, M. J. (2017). How is new technology changing job design? *IZA World of Labor*. https://doi.org/10.15185/izawol.344
- Gladden, M. E. (2016). Posthuman management: Creating effective organizations in an age of social robotics, ubiquitous AI, human augmentation, and virtual worlds. Defragmenter Media: USA.
- Grant, A. M. (2007). Relational job design and the motivation to make a prosocial difference. *Academy of Management Review*, 32(2), 393–417
- Grant, A. M. (2008). The significance of task significance: Job performance effects, relational mechanisms, and boundary conditions. *Journal of Applied Psychology*, 93(1), 108–124.
- Grogger, J., Ivandic, R., & Kirchmaier, T. (2020). Comparing conventional and machine-learning approaches to risk assessment in domestic abuse cases. *Journal of Empirical Legal Studies*, 18(1), 90–130.
- Hackman, J. R., & Oldham, G. R. (1975). Development of the job diagnostic survey. *Journal of Applied Psychology*, 60(2), 159–170.
- Hackman, J. R., & Oldham, G. R. (1976). Motivation through the design of work: Test of a theory. *Organizational Behavior and Human Performance*, 16(2), 250–279.
- Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30, 99–120.
- Hagras, H. (2018). Toward human-understandable, explainable AI. *Computer*, 51(9), 28–36.
- Halloran, L. & Andrews, J. (2018). Will you wait for the future to happen? Ernst and Young. https://www.ey.com/en_au/workforce/ will-you-shape-the-future-of-work-or-will-it-shape-you
- Hassabis, D., & Revell, T. (2021). With AI, you might unlock some of the secrets about how life works. New Scientist, 249(3315), 44–49.
- Hughes, J. (2014). A strategic opening for a basic income guarantee in the global crisis being created by AI, robots, desktop manufacturing and biomedicine. *Journal of Ethics and Emerging Technologies*, 24(1), 45–61.
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586.
- Jarrahi, M. H. (2019). In the age of the smart artificial intelligence: AI's dual capacities for automating and informating work. *Business Information Review*, 36(4), 178–187.
- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Manage*ment Annals, 14(1), 366–410.
- Langlois, R. N. (2003). Cognitive comparative advantage and the organization of work: Lessons from Herbert Simon's vision of the future. *Journal of Economic Psychology*, 24(2), 167–187.
- Leicht-Deobald, U., et al. (2019). The challenges of algorithm-based HR decision-making for personal integrity. *Journal of Business Ethics*, 160, 377–392.
- Lips-Wiersma, M., & Morris, L. (2009). Discriminating between 'meaningful work' and the 'management of meaning.' *Journal of Business Ethics*, 88(3), 491–511.
- Lips-Wiersma, M., & Wright, S. (2012). Measuring the meaning of meaningful work: Development and validation of the comprehensive meaningful work scale. *Group & Organization Management*, 37(5), 655–685.
- Lysova, E. I., Allan, B. A., Dik, B. J., Duffy, R. D., & Steger, M. F. (2019). Fostering meaningful work in organizations: A multi-level review and integration. *Journal of Vocational Behavior*, 110, 374–389.



- Martela, F., & Riekki, T. J. J. (2018). Autonomy, competence, relatedness, and beneficence: A multicultural comparison of the four pathways to meaningful work. Frontiers in Psychology, 9, 1157.
- Mazmanian, M., Orlikowski, W. J., & Yates, J. (2013). The autonomy paradox: The implications of mobile email devices for knowledge professionals. *Organization Science*, 24(5), 1337–1357.
- Michaelson, C., Pratt, M. G., Grant, A. M., & Dunn, C. P. (2014). Meaningful work: Connecting business ethics and organization studies. *Journal of Business Ethics*, 121, 77–90.
- Molitorisz, S. (2020). *Net privacy: How we can be free in an age of surveillance*. McGill-Queen's University Press: Canada.
- Naudé, W., & Dimitri, N. (2020). The race for an artificial general intelligence: Implications for public policy. AI & Society, 35, 367–379.
- Nussbaum, M. C. (2011). Creating capabilities: The human development approach. Harvard University Press: USA.
- Pardes, A. (2020, November). AI can run your work meetings now. Wired. https://www.wired.com/story/ai-can-run-work-meetings-now-headroom-clockwise/
- Parker, S. K., & Grote, G. (2022). Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Applied Psychology*, 71(4), 1171–1204.
- Pratt, M. G., & Ashforth, B. E. (2003). Fostering meaningfulness in working and at work. In K. Cameron, J. E. Dutton, & R. E. Quinn (Eds.), Positive organizational scholarship: Foundations of a new discipline (pp. 308–327). Berrett-Koehler: San Francisco.
- Pulse+IT. (2020). The San using AI to automate multidisciplinary team meetings. *Pulse+IT*. https://www.pulseitmagazine.com.au: 443/australian-ehealth/5558-the-san-using-ai-to-automate-multidisciplinary-team-meetings
- Ravenscraft, E. (25 November, 2021). What is the metaverse, exactly? Wired. Retrieved from: https://www.wired.com/story/what-is-the-metaverse/
- Roberts, P. (2020). Working smarter with data. *Australian Manufacturing Forum*. https://www.aumanufacturing.com.au/working-smart er-with-data-ai-gives-agriculture-the-competitive-edge
- Ryan, M., & Stahl, B. C. (2020). Artificial intelligence ethics guidelines for developers and users: Clarifying their content and normative implications. *Journal of Information, Communication and Ethics in Society*, 19(1), 61–86.
- Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Account*ability, and Transparency (pp. 59–68).
- Selenko, E., Bankins, S., Shoss, M., Warburton, J., & Restubog, S. L. D. (2022). Artificial intelligence and the future of work: A functional-identity perspective. *Current Directions in Psychologi*cal Science, 31(3), 272–279.
- Seppala, E., Rossomando, T., & Doty, J. R. (2013). Social connection and compassion: Important predictors of health and well-being. *Social Research*, 80(2), 411–430.
- Smids, J., Nyholm, S., & Berkers, H. (2020). Robots in the workplace: A threat to - or opportunity for - meaningful work? *Philosophy & Technology*, 33, 503–522.
- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*. https://doi.org/10.14763/2019.2.1410

- Symon, G., & Whiting, R. (2019). The sociomaterial negotiation of social entrepreneurs' meaningful work. *Journal of Management Studies*, 56(3), 655–684.
- Thaler, R., & Sunstein, C. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press: New Haven, CT
- Torres, P. (2019). The possibility and risks of artificial general intelligence. *Bulletin of the Atomic Scientists*, 75(3), 105–108.
- Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological Review*, 110(3), 403–421.
- Tubaro, P., Casilli, A. A., & Coville, M. (2020). The trainer, the verifier, the imitator: Three ways in which human platform workers support artificial intelligence. *Big Data & Society*, 7(1). https://doi.org/10.1177/2053951720919776
- Vallor, S. (2015). Moral deskilling and upskilling in a new machine age: Reflections on the ambiguous future of character. *Philosophy & Technology*, 28(1), 107–124.
- Walsh, T., Levy, N., Bell, G., Elliott, A., Maclaurin, J., Mareels, I., & Wood, Fiona. (2019). The effective and ethical development of artificial intelligence. ACOLA. https://acola.org/wp-content/uploads/2019/07/hs4_artificial-intelligence-report.pdf
- Wang, P. (2019). On defining artificial intelligence. *Journal of Artificial General Intelligence*, 10(2), 1–37.
- Webster, C., & Ivanov, S. (2020). Robotics, artificial intelligence, and the evolving nature of work. In: B. George & J. Paul (Eds.), Digital Transformation in Business and Society. Palgrave Macmillan, Cham.
- Wiesenfeld, B. M., Reyt, J.-N., Brockner, J., & Trope, Y. (2017). Construal level theory in organizational research. *Annual Review of Organizational Psychology and Organizational Behavior*, 4(1), 367–400.
- Wolf, S. (2010). Meaning in life and why it matters. Princeton University Press: New Jersey.
- World Economic Forum. (2018). *The future of jobs report*. Centre for the New Economy and Society: Geneva, Switzerland.
- Wright, D. (2011). A framework for the ethical impact assessment of information technology. *Ethics and Information Technology*, 13, 199–226.
- Wright, S. A., & Schultz, A. E. (2018). The rising tide of artificial intelligence and business automation: Developing an ethical framework. *Business Horizons*, 61(6), 823–832.
- Zhou, L., Paul, S., Demirkan, H., Yuan, L., Spohrer, J., Zhou, M., & Basu, J. (2021). Intelligence augmentation: Towards building human-machine symbiotic relationship. AIS Transactions on Human-Computer Interaction, 13(2), 243–264.
- Zuboff, S. (1988). In the age of the smart machine: The future of work and power. Basic Books: New York.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

