WORKER MANAGEMENT THROUGH AI

From technology development to the impacts on workers and their safety and health

1 Introduction and definition of research questions

Artificial intelligence (AI) is currently understood as one of the major threats to labour, although its use might in principle even ameliorate workers' safety and health if adequately designed, implemented, managed and regulated. Unlike conventional machines, the fundamental distinction of AI lies in its capacity to learn, coupled with the massive ability to store information through cloud computing. Learning has always served as a demarcation line between humans and machines, with the latter capacity being an inherent trait of human intelligence. If machines become capable of learning, as current AI algorithms allow, human-based activities in the work process have the potential to be replaced. However, concerns in terms of human tasks substitution by AI arise to a larger extent for so-called knowledge workers in fields such as media, communication, legal and accounting industries (Berg and Gmyrek, 2023).

As early as 2016, the Financial Times published an article titled 'When Your Boss is an Algorithm', recounting the story of the food delivery industry being entirely regulated through app-based processes. The capacity of AI to execute decision-making tasks has the potential to replace functions traditionally carried out by middle management. This transformation could also pose a challenge to the authority of higher-level supervisors, as well as result in a loss of relationships between workers and managers, which is an important mitigating factor of work-related stress. Over the past decade, academic research has extensively examined platform work and the changes in the work process resulting from the utilisation of digital systems.

In 'Your Boss is an Algorithm: Artificial Intelligence, Platform Work and Labour' by Aloisi and De Stefano in 2022, the authors assert that their book 'demonstrates that digital tools are more likely to replace managerial roles and intensify organisational processes in workplaces, rather than opening the way for mass job displacement'. Other scholars tend to concur with the significance of the managerial dimension beyond AI, considering it as one of the most relevant applications (Adams-Prassl, 2022).

However, the relevance of the transformation of the quality of work, as opposed to its impact on quantity, is a realm more commonly explored by sociologists and legal experts than economists (see Montobbio et al., 2022). One of the primary reasons why economists tend to focus more on quantity rather than quality is the absence of a measurement approach capable of quantifying several key factors: (i) the degree of technological penetration of such technologies, (ii) the specific scopes of their applications, and (iii) the functions and occupations most susceptible to exposure. If managerial applications are indeed as crucial as suggested, they should be demonstrable and quantifiable through specific indicators; otherwise, they may remain confined to isolated case studies and press releases.

This paper aims to bridge the existing gap and presents a measurement-based approach to analysing the diffusion and application of AI-based technologies designed for managerial functions. In line with the European Agency for Safety and Health at Work (EU-OSHA) definition (see EU-OSHA, 2022a, 2022c), this discussion paper introduces the term 'AI-based worker management' (AIWM) technologies and addresses the following questions: (i) Who are the inventors and key actors driving the development of these technologies? (ii) To what extent is the adoption of AIWM technologies spread across various sectors, going beyond specific case studies? (iii) What specific functions are these technologies intended to perform in relation to the workforce? (iv) What are the implications of these technologies for workers, particularly in terms of their safety and health?

To tackle these questions, we present a comprehensive mapping of AIWM technologies. This mapping involves in-depth analysis of patent data, gathering information about the key actors responsible for patent generation, examining the extent of sectoral application exposure, and analysing the textual content of patents to identify functions related to worker activities and tasks.

Our methodology is based on the work of Montobbio et al. (2022, 2023) and allows us to transition from a broad mapping of the knowledge base to a detailed analysis of sectoral penetration. It also has the potential to uncover specific occupations and worker tasks that are more targeted and exposed to AIWM technologies. This approach enables us to gauge the significance of AIWM in the application of such technologies right from the generation phase, as evident in patent texts. Furthermore, it helps us assess whether developers consider risks or opportunities for health and safety at work (Urzì Brancati and Curtarelli, 2021), even within the technological descriptions found in patent texts.

The paper is organised as follows. In section 1, we present the current state of the art and identify gaps in the literature. Section 2 outlines our methodology. Section 3 presents the results of our research. Section 4 discusses the implications for safety and health at work. Finally, we draw our conclusions.

2 State of the art and gaps in the literature

2.1 What are AIWM technologies?

In terms of technological definition, the literature often grapples with the challenge of precisely characterising what constitutes an AIWM technology and how it distinguishes itself from earlier software-based control technologies, which have been employed in various domains from production processes to labour management. An accepted definition describes AIWM technologies as ‘software algorithms that assume managerial functions, supported by surrounding institutional devices that facilitate the practical implementation of these algorithms’ (Lee et al., 2015, p. 1603). EU-OSHA (2002b, pp. 1-2) provides a more comprehensive definition, which reads as follows: ‘AI tools that are used to manage workers can be called AI-based worker management (AIWM), which refers to a worker management system that gathers data, often in real time, from the workspace, workers, the work they do and the (digital) tools they use for their work, which is then fed into an AI-based system that makes automated or semi-automated decisions or provides information for decision-makers (such as HR managers, employers, sometimes workers), on worker management related questions’. We will adhere to this definition henceforth.

2.2 What are their attributes and capabilities?

AIWM technologies, rooted in algorithmic structures, do not necessarily provide entirely new technological solutions. The use and development of algorithms and procedures can be traced back to earlier practices such as scientific management and time and motion studies (Dosi et al., 2021). However, what characterises AIWM is the degree of autonomous decision-making it exhibits in the execution of various functions and its role in complementing or substituting human decision-making activities. As outlined by Baiocco et al. (2022), what sets AIWM apart from conventional algorithm use is that these software systems have the capacity to make decisions, albeit to varying degrees (ranging from partial to complementary, and even involving a high level of automated decision-making). It is worth noting that fully autonomous decision-making is not yet deployed, and AIWM technologies appear to be more complementary or supportive rather than completely replacing human functions (Baiocco et al., 2022). Nevertheless, there is substantial heterogeneity across workplaces, and the extent of AIWM’s role depends on the specific business models and techno-organisational capabilities in place. According to EU-OSHA (2022c), AIWM encompasses a wide range of functions, including managerial ones, aimed at supporting worker productivity, and monitoring ones, related to directing, evaluating and enforcing discipline among workers. These two categories of functions are not mutually exclusive and are often simultaneously carried out by organisations.

2.3 What are their scopes of application?

Regarding the human–machine relationship, the literature reveals that AIWM technologies are employed to oversee (or automatically guide), assess and regulate work processes. They primarily assume managerial functions that are typically associated with white-collar and middle management roles, and these managerial functions directly influence mainly blue-collar workers heavily drawing on manual work and white-collar workers executing standardised and repetitive tasks (EU-OSHA, 2022c). From an organisational standpoint, these technologies appear to extend beyond the traditional Tayloristic role of control and supervision. Instead, they can displace low-end workers from their autonomy and areas of intervention (Moro et al., 2019; Cirillo et al., 2021). Simultaneously, AIWM technologies crystallise hierarchical structures, in effect transforming the ‘boss’ into an algorithm and disempowering workers. The capacity for (partial or semi-) autonomous decision-making, embedded within intelligent algorithms, can encompass a wide range of traditional firm activities. This includes regulating the pace and rhythm of production, managing the evolution of the internal labour market with regard to career progression and productivity incentives, making hiring and firing decisions (Adams-Prassl, 2022; EU-OSHA, 2022c; Urzì Brancati et al., 2022), and potentially even enabling worker profiling. Notably, according to EU-OSHA (2022c), AIWM technologies are more commonly adopted by organisations that maintain some form of worker representation. Although this finding may be influenced by a confounding effect related to the size of organisations — larger firms tend to have a greater need for AIWM practices and are also more likely to be unionised — it is important to note that there exists the potential for predictive applications in firms. Specifically, AI data collection on workers’ attitudes could enable predictive analysis, allowing organisations to detect anti-corporatist sentiment, misalignment with managerial decisions and disobedience to directives. This is made possible by predictive AI models that employ pattern recognition to categorise different types of workers based on their behaviour and attitudes. Indeed, this profiling capability could give rise to discriminatory practices against workers, for example against those who are more pro-union throughout their entire career life cycle (Bernhardt et al., 2023).

The reconfiguration of key managerial functions that AIWM technologies can undertake has been delineated, with specific reference to workforce management, encompassing the automation of direction, evaluation and discipline (Kellogg, 2020). To provide more detailed insight, EU-OSHA outlines three primary scopes of application (EU-OSHA, 2022c):

0. **Enhancing the efficiency and/or productivity of workers:** This scope encompasses functions such as task and shift allocation, as well as providing direction or guidance in the execution of various tasks.
1. **Enhancing the decision-making process:** Within this scope, functions include tracking employee performance, conducting evaluations, facilitating promotions, and even employing predictive models to anticipate worker decisions, such as leaving their current job.
2. **Enhancing workers’ health, safety and wellbeing:** This involves functions such as the identification, prevention and management of risky behaviours that may pose threats to workers’ health and safety.

2.4 What are the sectors of application?

Currently, the existing literature on AIWM technologies has primarily focused on the adoption processes and diffusion patterns, often relying on specific case studies. These studies have predominantly centred around platform work, particularly in the food delivery services sector (Wood, 2021). However, AIWM technologies have extended their reach beyond platforms to encompass various other sectors (EU-OSHA, 2022c). These include warehouses, supermarkets, call centres and healthcare facilities. In the case of warehouses and logistics, case studies, notably featuring Amazon, highlight the process of ‘human robotisation’ facilitated by digital technologies. This process entails increased control, standardisation and significant impacts on the work process (Massimo, 2020; Delfanti, 2021). Additionally, analyses have been conducted on the shipment processes within ports (Urzì Brancati et al., 2022). Beyond Amazon’s fulfilment centres, transportation on roads and last-mile delivery services have also experienced substantial transformations due to algorithmic management.

Call centres are a traditionally controlled sector, given the integration between communication systems, data on inputs (received calls), and outputs (fulfilment of scopes and number of sales) that have been recorded and stored in data format since the beginning of this industry. Additionally, call duration, the interval between calls and the number of breaks are all metrics that have been stored and computed by the managerial architecture, often referred to as an ‘electronic panopticon’ (Woodcock, 2022). AI-powered technologies applied to call centres, however, can do even more. For example, they can record biometric information like the tone of voice in order to predict the end result of the call, whether it will close with a sale or not. Indeed, selling is an intensely human-based activity, and telephone selling implies a high degree of persuasion and motivation from the seller (call centre operator) side. The use of AI-based technology to monitor workers’ emotions to predict stress in the execution of activities has also been recorded in some studies (Ball, 2021).

2.5 What are the implications for workers' safety and health?

The rallying cry of 'We are not robots', a pivotal slogan advocated by Amazon's labour force throughout the course of the workplace unionisation campaign,[1] serves as a succinct analogy that encapsulates the safety and health concerns emanating from AI-managed labour environments. According to the insights furnished by EU-OSHA (2022c), these concerns encompass multifaceted risks. , A frequently cited risk is work-related stress (EU-OSHA 2022b), which is exacerbated under the accelerated work paces often driven by automated algorithms. Instances such as traffic accidents involving delivery workers striving to meet automated delivery schedules underscore the ramifications of work-related stress induced by performance monitoring. Such stress is also a significant contributor to mental health issues. The second risk domain pertains to the management of a firm's internal labour market, encompassing aspects like rewards and career advancement, as well as disciplinary actions and layoffs based on information automatically generated by algorithms and with limited human intervention. Such poor HR management practices were found to generate feelings of job insecurity, anxiety and stress in workers (EU-OSHA, 2022b; Moore, 2029). Worker profiling also falls within this high-risk spectrum. A more comprehensive risk identification is available in EU-OSHA (2022b) that encompasses phenomena like increased work intensity and performance pressure, loss of workers' job control and autonomy, loss of social relationships with peers and managers, techno-stress, techno-anxiety, power imbalances and gamification, among others. Despite a substantial portion of these risks being well-recognised aspects inherent to occupational activities, the report underscores that organisations often fail to acknowledge AIWM as a risk factor. Additionally, risk assessments are frequently confined to the development and testing phases of AIWM, but AIWM systems tend to be neglected in the workplace risk assessment upon adoption in the workplace. Other risks mentioned in the literature are the dehumanisation, commodification and 'datafication' of human labour whereby, in the most extreme scenario, human labour might be relegated to the mere provision of data points (Mai, 2016; Negròn, 2021). This may force workers to "work like machines", leading to decreased autonomy, cognitive skills, creative thinking and critical thought, as well as individualisation, poorer worker's engagement and job satisfaction, which are known work-related psychosocial risks (EU-OSHA, 2022b). The absence of transparency regarding the utilisation and extent of AI-based monitoring presents a further set of risks (EU-OSHA, 2022b, 2022c). These range from privacy infringements to intrusions into the delicate balance between work and personal life, resulting in the potential alteration of individual personalities and the adoption of artificial behaviours to appease the algorithms, all of this resulting in an increase of stress and anxiety in workers (EU-OSHA, 2022b).

Opportunities for enhancing occupational safety and health (OSH) through AIWM do exist, primarily in the realm of pre-empting errors, resolving bottlenecks and averting accidents. In essence, the intellectual prowess underpinning AI algorithms can be harnessed for the betterment of OSH, issuing alerts and early warnings when critical risk factors are breached, scrutinising the OSH status within the workplace and offering recommendations for its enhancement. However, OSH-centric opportunities resulting from AIWM adoption have seen limited realisation, given that primary adoption objectives typically revolve around worker monitoring and productivity enhancement, as opposed to fostering worker wellbeing (EU-OSHA, 2022b). Furthermore, many of these opportunities, while promising, are inherently intrusive and all-encompassing. This includes initiatives such as mental health monitoring, digital counselling, and measures to boost worker engagement and satisfaction, all of which rely on individual data encompassing emotional, biometric and mental health attributes of workers. Consequently, these endeavours have the potential to engender negative cycles impacting worker OSH.

In the broader context, the acceptability of a panopticon-like system geared towards safeguarding worker health versus one centred on surveillance for control purposes remains a subject of ongoing debate.

2.6 Gaps in the literature

So far, the literature lacks large-scale studies. An exception is Urzì Brancati and Curtarelli (2022), who rely on the European Survey of Enterprises on New and Emerging Risks (ESENER) dataset, a company survey collecting data, among other information, on the adoption of management technologies among

[1] See: https://techcrunch.com/2018/11/23/amazon-warehouse-workers-in-europe-stage-we-are-not-robots-protests/

firms with more than five employees in EU countries. Sectors that mostly rely on such technologies are transportation and storage, manufacturing, and the electricity, gas and steam sector. However, beyond this, scholars are making inferential claims about the pervasive diffusion of such technologies, not only at the industry level but also at the workplace level. Adam Prassl (2022, p. 34) reports that 'algorithmic control is by no means limited to platform-based work. Employees across the socio-economic spectrum, from factories and offices to universities and professional services firms, are increasingly managed by, or with assistance from, algorithms. Start-ups and established software providers promise to automate employer functions across the life cycle of the employment relationship, from hiring and managing workers through to firing them'. The pandemic has accelerated the pace of adoption to control both the work carried out from home by white-collar workers and the pace of work of blue-collar workers (Adams-Prassl, 2022).

However, there are still gaps in the literature. Mast and Kresge (2022), focusing on the United States (US), emphasise that, in addition to the lack of large-scale evidence, the literature also lacks clear measures of adoption (prevalence of technology), has a lack of representativeness in case studies, an ambiguous definition of technologies and, finally, bias stemming from the source of private surveys from which the information is derived. Alongside these gaps, we argue that no study on technological developers and their heuristics has been put forward so far. The analysis of the heuristics of technological development allows us to address several gaps. Firstly, it helps in identifying who the technological developers are. To fill this gap, we will rely on patent data analysis. Secondly, it addresses the lack of verification of the scopes of application: do patented inventions genuinely exist to direct, control, evaluate and, alternatively, improve working conditions? Thirdly, it explores which functions, tasks and related occupations are most exposed to AIWM technologies. To address the latter question, we will construct a quantitative measure of worker exposure, going from tasks to occupations.

3 Methodology

The methodology we implement consists of several steps and makes use of multiple data sources. The latter include: PatentsView, from which we retrieve patents' full texts; CPC (v. 20230501) – Cooperative Patent Classification, which we use to identify patents within certain technological domains; Orbis IP, which connects patent documents to their corporate assignee(s) and includes industrial and geographic information thereof; and O*NET (v. 27.3),[2] which provides verbose description of human tasks and a means of aggregating them into occupations. Our methodological steps are depicted in the flowchart in Figure 1, while their extensive presentation and analytical details are presented in the Appendix.

Figure 1: Flowchart of our methodology

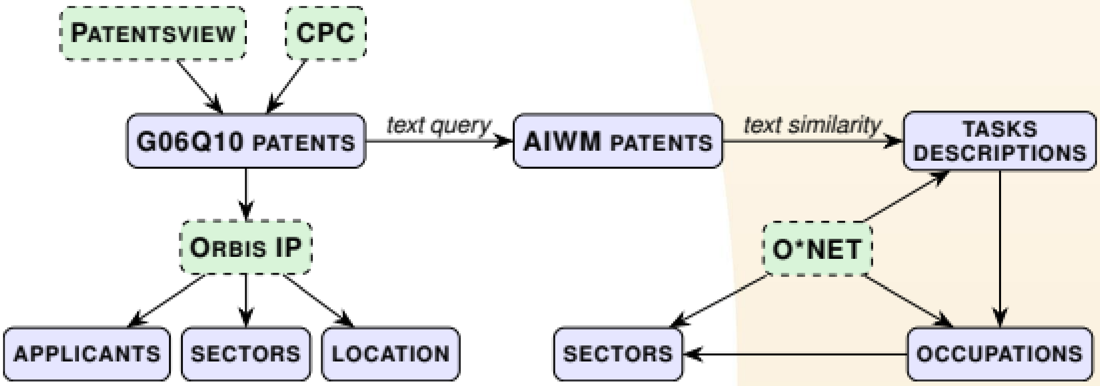

Source: Authors' elaboration

4 Results

In the following sections, we present our results categorised into three main domains. Firstly, we focus on the identification of the technology itself, its sectors of application, the inventors and their locations. After identifying the complete set of patent texts related to AIWM technologies using CPC code

[2] The Occupational Information Network (O*NET) is a free online database that contains standardised and occupation-specific descriptors on almost 1,000 occupations covering the entire United States economy.

definitions, we examine the scope of these patents (cf. step 1 of methodology). We narrow our focus to a subset of patents that explicitly report all three scopes of application identified earlier (cf. step 2). Secondly, we provide an overview of the results obtained from the similarity measure in terms of tasks. Finally, we report the results of the similarity measure in terms of occupations (cf. step 3).

Figure 2: Panel (a): number of G06Q10 patents per year (2002–2021); Panel (b): share of G06Q10 patents over all patents per year (2002–2021); Panel (c): number of G06Q10 patents per 4-digit NAICS sector of corporate holder; Panel (d): number of G06Q10 patents by country of corporate holder

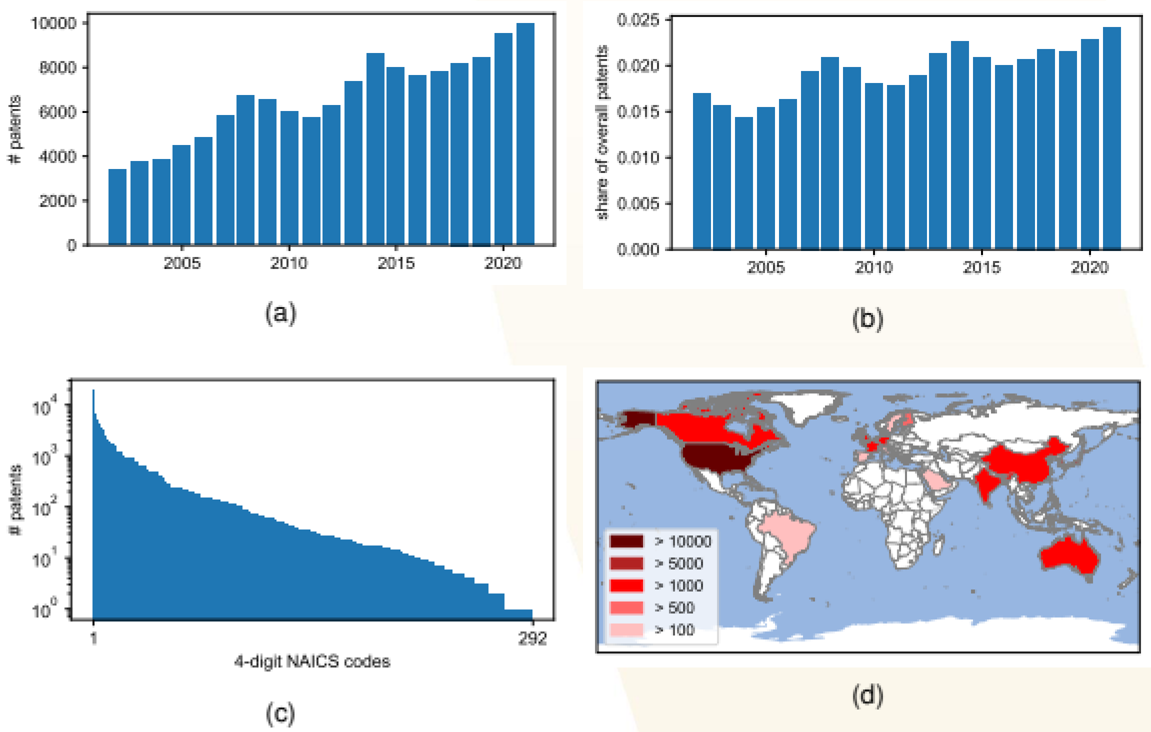

Source: Authors' elaboration

4.1 Technological mapping, sectors of application, inventors and their location

Figure 2 presents the results of technological mapping derived from the extraction of filed patents at the USPTO for the period 2002-2021 (i.e. all available data in whole years). When examining the CPC subclass 'Data processing systems or methods' (G06Q10), it is important to note that this subclass has been identified in the existing literature as a technological category where companies actively engage in patenting activities related to AIWM technologies. Panel (a) reveals a steep, upward trend spanning two decades. The initial numbers were around 4,000, gradually increasing to approximately 10,000 by the end of the period. In panel (b), we consider the time evolution of the share of G06Q10 patents relative to the total number of USPTO applications. This analysis allows us to control for the overall increase in patenting observed over the last 20 years. The figure confirms not only an increase in absolute numbers but also an upward trajectory in shares, indicating a 20-year growth rate of 42.8%. This substantial increase underscores the **genuine interest of innovation actors in patenting such technologies related to AIWM**.

Turning to the sectors of application for these technologies, we utilise a CPC-NAICS (North American Industry Classification System) concordance table to assess the concentration and diversification of the firms under analysis. Panel (c) provides evidence of a **widespread and dispersed distribution of sectors of application**, revealing a long right tail. In total, 292 sectors are represented. While certain dominant core sectors stand out (e.g. 'Computer Systems Design and Related Services' with 18,762 patents, 'Depository Credit Intermediation' with 6,191 patents, and 'Manufacturing and Reproducing Magnetic and Optical Media' with 5,119 patents), these innovations do not appear to originate exclusively from one sector. Notably, even sectors like 'Motor Vehicle Manufacturing' make the top 15 4-digit NAICS sectors with 1,692 patents, seemingly unrelated to their primary scope of innovation.

Finally, panel (d) illustrates the **geographical distribution of these inventors worldwide**, as a cumulative patent count. While there is a clear dominance of the US-Western region (86,301 patents), Germany (4,572) and Canada (3,416) also make substantial contributions. The Eastern region sees notable participation from Japan (9,792), South Korea (2,233), China (1,968), India (1,465) and Taiwan (939), all ranking among the top 15 countries. Additionally, Brazil and Saudi Arabia both have over 100 patents to their credit. These numbers are the result of the underlying firm-level patenting activity.

Table 1: Top applicants of 'Data processing systems or methods' (G06Q10) patents (2001-2022)

Rank	# patents	Company name	Country	% portfolio
1	7906	IBM	US	1.62%
2	4821	Microsoft	US	2.77%
3	2515	SAP	DE	8.33%
4	1581	Google	US	1.16%
5	1480	JPMorgan	US	1.72%
6	1207	Bank of America	US	1.63%
7	1179	Oracle	US	3.98%
8	907	Walmart	US	11.41%
9	874	Fujitsu	JP	0.17%
10	773	Bank of America N.A.	US	1.79%
11	753	Hitachi	JP	0.09%
12	739	Blackberry	CA	1.10%
13	695	Toyota	JP	0.12%
14	639	Samsung	KR	0.06%
15	608	HP	US	0.52%

Source: Authors' elaboration

Table 1 provides a breakdown of the top applicants for G06Q10 patents over the entire period. The data were sourced from Orbis IP, and it is important to note that the overall disambiguation of data may not be perfect. US-based companies dominate the list, with IBM leading with 7,906 patents, followed by Microsoft with 4,821 patents. Additionally, Google (1,581), JPMorgan (1,480) and Bank of America (1,207) are among the top 10 inventors throughout. When examining the most specialised firms in patenting activity, particularly those with substantial patent portfolios, we consider the share of 'Data processing systems or methods' (G06Q10) patents relative to their entire patent portfolio. Among the largest inventors, Walmart stands out with 11.47% of all G06Q10 patents, totalling 907 patents, making it the most specialised in this category. Following closely is SAP, with a share of 8.33% out of a total of 2,515 patents in G06Q10. The third most specialised company is Oracle, with a 3.98% share out of a total of 1,179 'Data processing systems or methods' (G06Q10) patents. The share indicator provides insights into the intensity of innovative activities in the specific technology of interest. While firms like SAP and Oracle might be more expected in this context, the presence of large traditional retailers like Walmart in this list is less anticipated. Once again, this evidence underscores the **wide-ranging application of these technologies**.

Table 2: Breakdown of CPC group G06Q50

Rank	# patents	CPC group	CPC title	% G06Q10
1	17399	G06Q50/01	Social networking	35.21%
2	5342	G06Q50/30	Transportation; Communications	10.81%
3	5105	G06Q50/06	Electricity, gas or water supply	10.33%
4	2624	G06Q50/28	Logistics, e.g. warehousing, loading, distribution or shipping	5.31%
5	2219	G06Q50/12	Hotels or restaurants	4.49%
6	2173	G06Q50/265	Personal security, identity or safety	4.40%
7	1896	G06Q50/26	Government or public services	3.84%
8	1766	G06Q50/10	Services	3.57%
9	1599	G06Q50/02	Agriculture; Fishing; Mining	3.24%
10	1583	G06Q50/04	Manufacturing	3.20%
11	1499	G06Q50/184	Intellectual property management	3.03%
12	1494	G06Q50/18	Legal services; Handling legal documents	3.02%
13	1481	G06Q50/16	Real estate	3.00%
14	1343	G06Q50/22	Social work	2.72%
15	1018	G06Q50/14	Travel agencies	2.06%

Source: Authors' elaboration

One initial insight into the scope of these technologies can be gleaned from the CPC classification G06Q50 ('Data processing systems or methods … specially adapted for specific business sectors') itself, which, unlike G06Q10 ('Data processing systems or methods'), further categorises applications into specific sectors. Consequently, when we delve into finer detail and calculate the ranking of sectoral applications for the second decade (2012-2022), we discover the following **top sectors:**

- **'Social networking'**, represented by G06Q50/01 with 17,399 patents, constituting approximately a third of the total G06Q50 patents;
- '**Transportation and Communications'**, denoted as G06Q50/30, with 5,342 patents;
- '**Electricity, gas, or water supply'**, identified as G06Q50/06, with 51,055 patents, both sectors accounting for roughly 10% of the total G06Q50; and
- '**Logistics, including warehousing, loading, distribution, or shipping'** (G06Q50/28), with 2,624 patents, representing 5% of the total.

Notably, these sectors of activity to which these technologies are applied heavily rely on functions such as communication, information processing, data analysis and human interaction. In the subsequent sections, we will conduct a more detailed textual content analysis to gain a finer-grained understanding of their scope.

When we examine the patterns of co-occurrence of group G06Q10 with other CPC codes, which indicate complementary domains of application, we observe interesting trends. In the 2012-2022 decade, the highest co-occurrence is with 'Machine learning' (G06N20/00). In contrast, during the first decade (2001-2011), the highest co-occurrence is with 'Commerce, e.g., shopping or e-commerce-Marketing' (G06Q30/02). This finding further substantiates the set of identified patents, as it demonstrates that they **encompass not only specific administrative and managerial functions but are also clearly enhanced by machine learning algorithms that support their operations**. With this reliable set of patents incorporating AIWM, we can now proceed to the subsequent steps of our analysis.

4.2 Full-text analysis: identification of patent scopes of application

To gain deeper insights into the specific scopes of the identified set of 'Data processing systems or methods' (G06Q10) patents, we conduct a textual analysis on the patent text corpora, enabling us to refine our search. Commencing with the set of G06Q10 patents, we execute a structured text query designed to ascertain whether our identified patents indeed encompass the scopes of application as defined above.

In the following section, we present excerpts from the 4,340 patents that have been identified, along the three AIWM heuristics:

1. Increase efficiency and/or productivity of workers:

   '… firms may also monitor information relating to the production rates and the performance of individual production workers, production lines, work centres, production teams and pieces of production equipment …' (US20010041995A1)

2. Improve the decision-making process:

   '… the efficiency management system is configured as a "personnel work management system" adapted for evaluating employee or staff efficiencies based on data indicative of activity and time …' (US20140188533A1)

3. Improve workers' health, safety and/or wellbeing:

   '… the dynamic inputs may include a human resources input and a mobile resources input that may be used to calculate, manage, and mitigate risk that may be associated particularly with humans being detected in certain possibly hazardous and/or restricted locations, spaces, or areas of the industrial facility.' (US20140156584A1)

It is important to note that the presence of the queried text was not initially anticipated, as the technical descriptions of a technology do not always explicitly refer to the human–machine relationship. However, **through our analysis of patent texts, we can infer not only the scopes but also the attributes and capabilities of this technology. These range from the classification of information to enhanced managerial decision-making and prediction.**

Figure 3: Distribution of cosine similarity measure

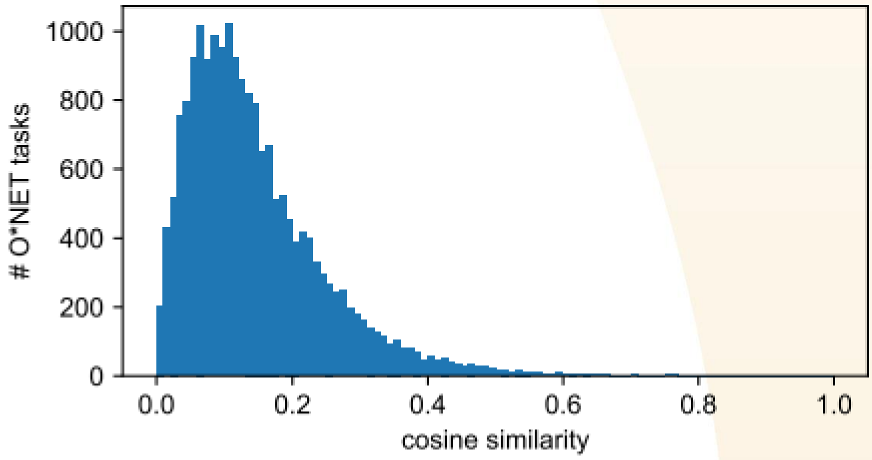

Source: Authors' elaboration

4.3 Text-similarity measure: identification of worker tasks and occupations

Having identified a subset of 4,340 patents reporting specific scopes of applications related to increasing efficiency, monitoring, improving decision making and safety at work, we now present the results of the text similarity measure, the detailed construction of which is provided in the Appendix. This measure quantifies the degree of similarity between two text corpora: one derived from the set of identified patents and the other from the descriptions of human functions and tasks performed by workers, using the O*NET dataset. This allows us to specifically examine the types of human functions most closely related with these technologies.

In Figure 3, we present the distribution of the cosine similarity measure between the patent texts and the O*NET tasks. The distribution exhibits significant skewness, with high similarity values approaching one being a rare occurrence. By focusing on this subset of tasks, we gain a **better understanding of the functions that these technological applications either substitute, complement, enhance or support**.

Figure 4: Word cloud of tasks descriptions, rescaled by cosine similarity

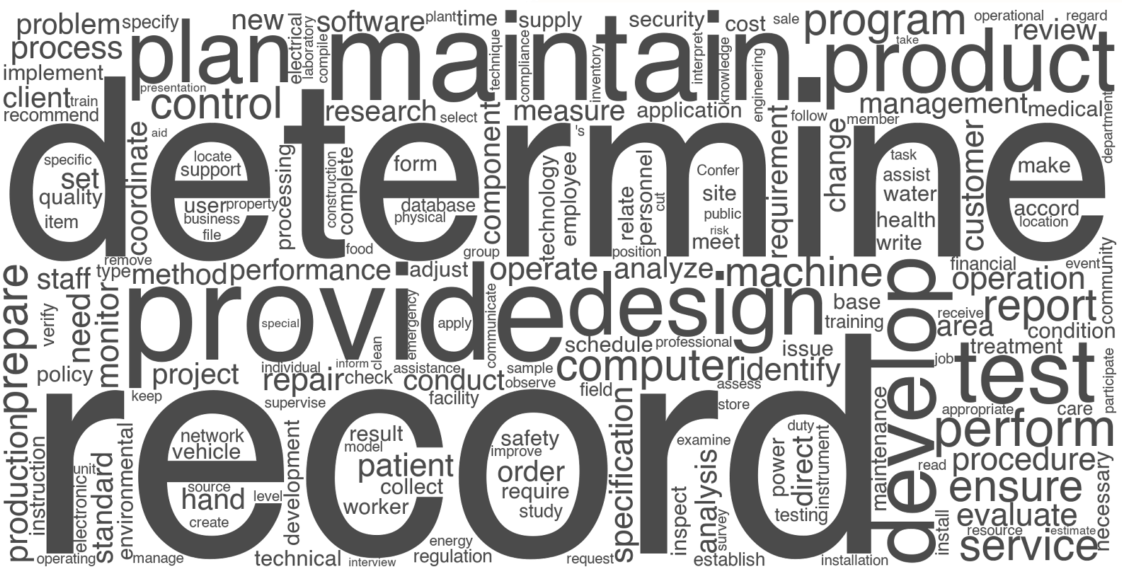

Source: Authors' elaboration

Figure 4 presents a word cloud comprising relevant terms found in task descriptions. The frequency of words in each task has been rescaled based on the underlying cosine similarity measure. English stop words have been excluded (cf. Appendix). Additionally, generic terms like 'system', 'activity' and 'equipment', which are applicable regardless of the specific task, have also been omitted.

We reiterate that the similarity measure is constructed based on the entire corpus of selected patent texts, which is not limited to the sentences identified by the text query (see Figure A 1 in Appendix), and the task descriptions in the O*NET. The identification of the most relevant words is determined by a similarity measure that takes into account the degree of overlap between these two text corpora, while discounting general overlapping words and assigning more weight to specific, characteristic words. Consequently, the cosine similarity matrix tends to overemphasise rare yet pertinent information while downplaying common, less specific information.

According to this measure, the top emerging words in terms of similarity are 'determine', 'provide', 'plan', 'maintain', 'develop' and 'test'. These words represent the most frequently overlapping functions between those executed by AIWM technologies and human activities. This result initially validates the identification procedure. Notably, although we began with a defined set of patents based on specific semantic searches, the most frequent functions are not necessarily the ones explicitly sought after; they instead represent a broader and more abstract meta-domain of activities, transcending strict supervision and control, and encompassing decision-making functions based on autonomy and independent judgment, thereby enhancing intelligent operations characteristic of human capabilities. **The most frequently overlapping functions seem to fall within the scope of 'Improve the decision-making process'.**

Words such as 'record', 'performance', 'monitor', 'evaluate', 'inspect', 'schedule' and 'standard' provide insight into the realm of monitoring performance and evaluating work processes. This group of words is more **closely related to the scope of 'Increase efficiency and/or productivity of workers'**, which is more directed at controlling and self-defining production standards and error tracking. While some overlap with the semantic search exists (e.g. 'performance'), these words constitute a complementary dictionary of functions to those explicitly sought, representing a larger, semantically related set of functions.

Finally, words such as 'patient', 'safety', 'medical' and 'environmental' are associated with the scope of 'Enhancing workers' health, safety, and well-being' and offer insight into the **potential applicability in terms of enhancing the health and safety dimension. However, these words and related functions are significantly less prevalent.**

Table 3: Two-digit SOC occupations by level of exposure (normalised between 0 and 1)

Rank	2-digit SOC	Norm. exposure
1	Production Occupations	1.000
2	Architecture and Engineering Occupations	0.870
3	Life, Physical, and Social Science Occupations	0.761
4	Office and Administrative Support Occupations	0.758
5	Management Occupations	0.671
6	Healthcare Practitioners and Technical Occupations	0.651
7	Computer and Mathematical Occupations	0.622
8	Business and Financial Operations Occupations	0.561
9	Installation, Maintenance, and Repair Occupations	0.452
10	Transportation and Material Moving Occupations	0.412
11	Construction and Extraction Occupations	0.293
12	Arts, Design, Entertainment, Sports, and Media Occupations	0.276
13	Educational Instruction and Library Occupations	0.222
14	Sales and Related Occupations	0.219
15	Personal Care and Service Occupations	0.178
16	Food Preparation and Serving Related Occupations	0.160
17	Healthcare Support Occupations	0.126
18	Protective Service Occupations	0.125
19	Community and Social Service Occupations	0.120
20	Farming, Fishing, and Forestry Occupations	0.019
21	Building and Grounds Cleaning and Maintenance Occupations	0.011
22	Legal Occupations	0.000

Source: Authors' elaboration

Moving beyond functions and activities, it is pertinent to inquire about the occupations most closely associated with AIWM technologies. Table 3 presents the relative ranking of two-digit occupations in terms of similarity. The top-related occupations are within production processes, a result in line with expectations, as **AIWM systems generally operate in standardised environments, necessitating the definition of precise procedures, operations and workflows** typical of the manufacturing sector.

However, from the second to the ninth most similar occupations, those belonging to the upper echelon of the occupational structure emerge. These are **engineering, technical and specialised professions that perform autonomous decision-making functions**, such as 'determine', 'conduct', 'analyse' and 'record' (cf. Figure 4). Consequently, a high level of similarity is primarily observed with managerial activities. An illustrative application of AIWM technologies, designed to augment human decision-making, is witnessed in Human Resource (HR) departments where the practice of people analytics is on the rise. According to Moore (2019), transnational and large-scale corporations are progressively integrating AI-based algorithms into HR activities, primarily aimed at addressing the intricacies associated with managing the so-called people problem. The occupations responsible for executing HR functions are predominantly rooted in the domains of social science and management. Additionally, AIWM technologies extend their influence to fields beyond HR, where engineering and life science occupations also find their roles complemented by these advanced tools. For instance, consider the case of MUSE,[3] a cutting-edge smart headband that leverages brain–computer interface technology. This innovation can assist employers in the execution of wellness programmes, as expounded by Gonfalonieri (2020).

To which occupations are these managerial (more or less artificial) functions directed? The lower echelon of the occupational structure emerges from the ninth rank onwards, including primarily manual

[3] See: https://choosemuse.com/

occupations in activities such as 'Installation, Maintenance, and Repair', 'Transportation, and Material Moving', as well as 'Construction and Extraction'. Standardised office-based occupations in 'Sales' and 'Educational Instruction and Library' also fall into this category. Indeed, some of these occupations are more common in logistics and call centres, which are those with the most targeted workplaces for AIWM applications (EU-OSHA, 2019).

'Healthcare', a sector of particular interest[4] ranks sixth in similarity, while 'Personal Care and Services', which involve less dependent relationships and greater autonomy, rank 15th. In general, setting aside production workers in manufacturing, Table 3 reflects a **clear hierarchical nature of the relationship with AIWM technologies, which are closer to the execution of functions typical of autonomous managerial occupations and are aimed at guiding and directing non-autonomous, dependent and subordinated occupations.**

The specific occupations that benefit from AIWM-enabled management align with broader findings in the literature, which have often focused on sectors rather than individual job roles. These findings consistently identify healthcare, logistics and manufacturing as the industries with the highest adoption rates (EU-OSHA, 2022b, p. 26). Our results also underscore the notable similarity between standardised white-collar occupations, such as 'Office and Administrative Support Occupations', and 'Business and Financial Operations Occupations'. Consequently, **our research not only corroborates previous evidence but also provides a more detailed perspective on the extent to which AIWM technologies permeate various occupational categories**.

The cosine similarity measures among occupations (Figure 3) exhibit a gradual decline from 1 to 0 without abrupt discontinuities. However, this is not the case when occupations are categorised into sectors, as shown in Table 4. Each sector presents a unique occupational composition, and as a result, when aggregated into sectors, the similarity distances vary, ranging from 1 to 0.543 (in relative terms) between the first and second most exposed sectors of application. This greater distance primarily arises from the inherently more homogeneous nature of manufacturing sectors versus less homogeneous service sectors in terms of occupational composition. Consequently, production occupations make manufacturing the most exposed sector. That said, services also appear notably exposed. Within the top five sectors, 'Health care and social assistance', 'Education services', 'Public administration', and 'Transportation and warehousing' all feature prominently. The presence of these applications across various service sectors is further substantiated by case study evidence, which reports the utilisation of software implementing forms of algorithmic management in a diverse range of service industries such as janitorial and security services, restaurants, hotels, healthcare and the public sector. This complements the previously mentioned call centre and transport industries (Bernhardt et al., 2021).

[4] The human health and social work activities sector is one of the largest in Europe, employing around 11% of workers in the European Union according to 2020 figures from Eurostat (see at https://osha.europa.eu/en/publications/human-health-and-social-work-activities-evidence-european-survey-enterprises-new-and-emerging-risks-esener). It represents the most rapidly growing segment of demanded jobs according to the U.S. Bureau of Labor Statistics.

Table 4: Two-digit NAICS sectors by level of exposure (normalised between 0 and 1)

Rank	2-digit NAICS	Norm. exposure
1	Manufacturing	1.000
2	Health care and social assistance	0.543
3	Public administration	0.390
4	Education services	0.383
5	Transportation and warehousing	0.284
6	Professional, Scientific and Technical Services	0.205
7	Other services, except public administration	0.187
8	Retail trade	0.170
9	Utilities	0.166
10	Construction	0.164
11	Information	0.157
12	Finance and insurance	0.137
13	Accommodation and food services	0.136
14	Arts, entertainment, and recreation	0.124
15	Administrative and support and waste management and remediation services	0.095
16	Wholesale trade	0.072
17	Mining	0.053
18	Real estate and rental and leasing	0.034
18	Agriculture, forestry, fishing and hunting	0.034
20	Management of Companies and Enterprises	0.001

Source: Authors' elaboration

5 Implications for OSH

The analysis conducted has revealed the remarkable **pervasive nature of these technologies and their potential to either combine with or substitute for many human functions**. Although potentially disruptive in their effects, technologies themselves are inherently neither good nor bad. The non-neutrality of their effects depends on human decisions regarding how to adapt and use these tools.

Interviews conducted in the second half of 2023 as part of this project with two chief OSH managers in multinational companies based in Europe that develop and use AIWM technologies allowed to collect qualitative information based on their expert point of view on the extent of these technologies' pervasiveness in workplaces. Firstly, they perceive room for adaptation and change management through the **influence of the corporate culture** within their respective firms. Secondly, they identify potential **benefits in many workplaces for improving OSH**. For example, workplace AI can support functions that help individuals find solutions to health problems, prevent burnout and promote a better work–life balance, and therefore implement form of digital counselling and use AIWM to increase OSH at the workplace level. Thirdly, they mentioned the **significance of legal provisions**, such as the EU General Data Protection Regulation (GDPR), that enforces similar legislative restrictions on non-EU firms. Lastly, these managers interviewed reckoned that **AIWM technologies are currently in their infancy of application, without specific targeted human functions and occupations envisaged**, and are primarily relevant in industries such as transport and delivery, while having less impact on working conditions in higher-skilled jobs. The interviewees perceive the possibility of **shaping and directing the adoption patterns of these technologies at the workplace level** through corporate institutions, boards and human agency capable of mitigating potential distortions. In general, **AIWM technologies are viewed as neutral**.

The experts interviewed for this project as well as existing research, share the idea that there is no predetermined outcome, and these technologies might also be employed to improve working conditions (EU-OSHA, 2022b). As stated by Bernhardt et al (2023, p.5):

> 'the outcomes of employer adoption of emerging technologies are not predetermined. The social and institutional context shapes employer decision-making about new technologies, leading to variation in adoption and worker impacts.'

Indeed, the literature reports certain benefits of AIWM systems for employers as well as for OSH, and some providers of AIWM such as HR analytics software advertise OSH benefits among their selling arguments. However, **even with good intentions, such systems still might negatively impact workers' health, safety and wellbeing** (EU-OSHA, 2022b). Bernhardt *et* al. (2021), in a report released by the UC Berkeley Labor Center[5] reviewing examples of practical applications of AIWM technologies in the workplace, present a similar perspective to EU-OSHA's research (2022b) when it comes to identifying risks workers' health. These include work intensification, reduced autonomy and job control, deskilling, discrimination, fear of job loss, lower wages, reduced economic mobility, contingent work, suppression of the right to organise, and a decline privacy and dignity, which can all impact mostly on workers' mental health.

In contrast to the perspectives of the two experts interviewed, our extensive data analysis reveals a **prevalence of OSH risks rather than opportunities when examining patent texts and their similarity to worker tasks**. Notably, the majority of the retrieved patents predominantly focus on scopes of application aimed at **enhancing efficiency and improving decision-making processes**. The most frequently occurring terms within the word cloud (cf. Figure 4) highlight a collection of functions performed by AIWM technologies, primarily centred on recording, directing and assessing. These technological scopes are unmistakably oriented towards **worker monitoring, rating and evaluation**, which ultimately facilitate digital and distributed workforce surveillance.

The ramifications of AI-based monitoring for worker safety and health are profound, encompassing a spectrum of **both mental and physical work-related illnesses**. These include techno-stress, techno-anxiety and work-related stress, as well as physical injuries and accidents. These risks are intrinsically associated with the objective of improving productivity. However, more nuanced and emerging forms of risk are also on the horizon. These encompass phenomena like gamification, privacy infringements, and the broader concerns of dehumanisation and 'datafication' of workers and workplaces, as documented by EU-OSHA (2022b). These risks are closely tied to the overarching goal of enhancing managerial decision-making processes, which is made possible through extensive data collection, AI processing for pattern recognition, and the subsequent application of these insights to tasks such as shift organisation, performance evaluation and the prediction of anti-corporate behaviours.

Overall, AIWM technologies are progressively tilting the power balance in favour of managerial prerogatives at the expense of worker empowerment. Simultaneously, these technologies pose a threat to workers' abilities to build coalitions, voice collective grievances and engage in unionisation efforts, as noted by Negròn (2021). Their implementation may also jeopardise workers´ (and their representatives´) rights to know the purpose of AIWM and to have access to the data. However, at the time of writing this discussion paper, the EU Institutions are negotiating regulatory provisions to regulate worker´s rights to access data collected, processed and used through algorithmic management through the proposed directive to improve working conditions for platform workers[6] and Artificial Intelligence Act for trustworthy AI (AI Act)[7]. Some Member States, such as Spain[8], already paved the way in this direction. Italy even took an attempt to go beyond granting data access to workers and their representatives by also granting access to labour inspectors through the introduction of the Legislative

[5] This refers to research in the US, however, the findings are equally relevant to workplaces in the EU even though the OSH regulatory context is different in the US.

[6] https://www.consilium.europa.eu/en/press/press-releases/2023/12/13/rights-for-platform-workers-council-and-parliament-strike-deal/

[7] https://www.europarl.europa.eu/news/en/press-room/20231206IPR15699/artificial-intelligence-act-deal-on-comprehensive-rules-for-trustworthy-ai

[8] In Spain, the Spanish Royal Decree-Law 9/20213 (also known as the Riders' law) created the obligation of information and consultation of workers' representatives not only to digital labour platforms but also to all organisations using algorithmic or AI-based systems which could affect working conditions. See: https://www.mites.gob.es/ficheros/ministerio/inicio_destacados/Guia_Algoritmos_EN.pdf

Decree104/2022 that transposes the EU Directive on Transparent and Predictable Working Conditions[9] (Pesole, 2023). Beside the legislative developments attempting to regulate algorithms specifically, it is important to emphasise that the OSH Framework Directive[10], although generic in its scope, implicitly applies to AIWM, meaning that employers have the obligation to inform workers about health and safety risks associated with AIWM and to engage in consultation with workers. However, the information asymmetries between employers and workers that AIWM may create have the potential to impede a proper worker consultation in practice.

The mere existence of patents that incorporate worker monitoring functions, which are on the rise over time, filed by a diverse array of companies ranging from big tech firms to financial corporations and retail giants, underscores the pervasive corporate interest in technological solutions that facilitate worker management through AI-based algorithms. This, in turn, signals the widespread presence of OSH risks posed to workers by these technologies.

We fully agree that **workplaces remain the crucial context for understanding whether risks or improvements will result from the adoption of AIWM technologies**. Despite our large-scale analysis, we do not provide detailed insights into the actual patterns of adoption at the workplace level. Therefore, we consider firm-level surveys conducted at the EU level to be a vital source for understanding the practical application of such technologies and to determine the extent to which human-based management of change will be implemented. Additionally, while discrimination and anti-union practices may be less applicable in more regulated EU labour markets, they are still risks that need to be prevented through legislation.

6 Conclusions

We have undertaken an extensive large-scale analysis of AIWM technologies to discern their technological landscape, functionalities and scopes. In a noteworthy departure from existing literature, we have executed a systematic general-to-specific analytical approach, commencing with a focus on the technologies themselves and subsequently extending our investigation to encompass their impact on the workforce, tasks and functions. Our findings unveil several salient trends:

1. **Continuous Technological Advancement:** These technologies have demonstrated a consistent upward trajectory, both in terms of their absolute numbers and their expanding market presence.
2. **Identification of Leading Patent-Holders:** We have successfully identified prominent patent-holding entities, encompassing not only industry leaders within the digital technology sector but also significant players in complementary industries such as retail trade, banking and automotive.
3. **Global Distribution of Key Players:** Our geographical analysis has pinpointed the prevalence of US-based firms, alongside well-established entities in the EU and Japan, as well as emerging contenders in countries like China and South Korea.
4. **Diverse Scopes of Application:** We have identified diverse scopes for the application of AIWM technologies, including the improvement of worker efficiency, facilitation of decision-making processes, and enhancement of health and safety measures.
5. **Precise Identification of Tasks and Occupations:** Employing a quantitative textual analysis, we have meticulously identified specific tasks and occupational categories directly affected by AIWM technologies.

In the realm of tasks, our research has highlighted activities closely tied to decision-making processes, such as determining, providing, maintaining and recording. These tasks are intrinsically linked to human cognitive abilities and play a pivotal role in defining workflow and production processes. With regard to occupations, our study corroborates that those most significantly impacted are primarily situated within the manufacturing sector. However, a distinctive pattern emerges when examining the distribution of occupations. **Managerial-related roles appear to be the principal target of AIWM technologies, while, conversely, dependent and subordinated occupations appear to be the primary focus of these technological advancements.** The far-reaching influence of these technologies is further

[9] https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019L1152
[10] See: https://osha.europa.eu/en/legislation/directives/the-osh-framework-directive/1

underscored by their extensive penetration across a wide spectrum of sectors, encompassing not only manufacturing but also a diverse array of service-based activities.

It is essential to acknowledge the primary limitation of our research, which predominantly revolves around the absence of empirical evidence regarding workplace adoption. To address this limitation, we recommend the exploration of large-scale workplace surveys as a promising avenue for gathering valuable insights into the practical implications of AIWM technologies. Furthermore, we recognise the need for in-depth investigations into the realm of discriminatory practices, particularly concerning hiring, termination and worker profiling. This area presents a compelling avenue for future research, with a particular focus on its implications for labour unions and worker behaviours.

In conclusion, our analysis contributes to the pool of research on the adverse impact that the management of workers through algorithms or AI can have on wellbeing, encompassing safety at work and mental and physical health. While we have delineated potential opportunities of AIWM for OSH, our empirical search has revealed **a limited commitment to innovative endeavours towards enhancing worker's OSH through AIWM technologies**. We have retrieved only a marginal fraction of patents aimed at improving human wellbeing, underscoring the dearth of corporate interest in employing AIWM technologies to improve OSH. Furthermore, AIWM technologies, even when employed to enhance OSH, may introduce intrusive and controversial side effects due to their comprehensive capacity to collect and analyse sensitive worker data encompassing health status, mental and physical conditions, psychological profiles, social attitudes and even brain activity, as articulated by Gonfalonieri (2020).

Additionally, the lack of transparency and information asymmetry concerning the use of AI devices can hinder worker agency and potentially heighten the tacit acceptance of these technologies and the related OSH risks. Hence, it is advisable that research endeavours be initiated to categorise the degree of pervasiveness, incidence and severity associated with the full spectrum of OSH-related risks. Moreover, we advocate for clear and unambiguous provision of information. It is imperative that social dialogue actors and policymakers mandate transparent disclosure when an AI-empowered algorithm assumes managerial responsibilities in a given workplace (De Stefano and Taes, 2023). This transparency ensures that workers are fully informed about the AI's role and its potential implications for their OSH. The proposed directive to improve working conditions for platform workers[11] and Artificial Intelligence Act for trustworthy AI (AI Act)[12], which the EU Institutions are negotiating at the time of writing, include provisions to regulate the transparency of algorithms and worker´s rights to access data collected, processed and used through algorithmic management, as this is already the case some Member States[13].

[11] https://www.consilium.europa.eu/en/press/press-releases/2023/12/13/rights-for-platform-workers-council-and-parliament-strike-deal/
[12] https://www.europarl.europa.eu/news/en/press-room/20231206IPR15699/artificial-intelligence-act-deal-on-comprehensive-rules-for-trustworthy-ai
[13] See examples of Spain and Italy mentioned previously.

Appendix

This Appendix presents the details of the methodology and serves as a technical summary of the methodological details of our analysis.

Identification and mapping of the technology of interest

We employ a general-to-specific methodology, commencing with a clear definition of the technology and subsequently examining applicants, locations and sectors of exposure. The analysis begins with the entire set of 6,952,896 patent applications published by the United States Patent and Trademark Office (USPTO) between 15 March 2001 and 30 June 2022. Within this dataset, we initially identify a broad set of 138,288 patents, expected to contain AIWM patents, among others. We achieve this by utilising patent hierarchical technological classification codes. Specifically, we focus on CPC (Cooperative Patent Classification) group G06Q10, which pertains to 'Data processing systems or methods' (cf. subclass G06Q) 'specially adapted for Administration; Management', and lower-digit codes within its subgroup. These codes encompass information and communication technologies (ICT) specially adapted for the management of organisations or their employees. This includes automation of office work using computers and time management in an enterprise environment, such as monitoring billable hours and working time accounting for employees.[14] These codes were selected because they represent patents containing relevant keywords, such as 'algorithmic management', and are prevalent in patents filed by the focus companies listed in the Coworker report, a leading source on this matter.[15] Simultaneously, to gain insights into the technological and industrial diversity of these patents, we examine the co-occurrence of CPC codes of group G06Q10 with CPC codes from other groups or classes. It's worth noting that a patent can be assigned multiple CPC codes when deemed appropriate by the underlying patent examiner, which is usually the case. In one instance, we employ another CPC group under subclass G06Q, namely G06Q50, which encompasses 'Data processing systems or methods … specially adapted for specific business sectors'. According to CPC rules, systems or methods in group 'Data processing systems or methods' (G06Q10) that are specially adapted for a specific business sector must also be classified in group 'Data processing systems or methods … specially adapted for specific business sectors' (G06Q50) when the special adaptation is deemed novel and non-obvious. Lower digits of G06Q50 refer to the underlying sectors, including Agriculture/Fishing/Mining (G06Q50/02), Manufacturing (G06Q50/04), and a selection of services, such as Legal (G06Q50/18), Education (G06Q50/20) and Logistics (G06Q50/28). The analysis in this step enables us to identify the actors responsible for relevant technological innovations, assisting us in detecting who they are and where they are located. Furthermore, it allows us to study the sectoral penetration in terms of the scope of application of algorithmic management technologies.

Identification of scopes of technology and impact upon the workforce

This methodological step draws significant inspiration from Montobbio et al. (2022). AIWM patents are a subset of G06Q10 patents, identified using a multi-word co-occurrence query at the sentence level. To achieve this, we go through several preprocessing steps. First, we break down each G06Q10 patent's description into a list of sentences. Second, we similarly break down each sentence into a list of words. Third, we filter out a standard set of 182 common English stop-words, such as 'a', 'the' and 'if', which do not provide valuable information for our analysis. Finally, we reduce each word in each sentence to its root form using a lemmatisation algorithm.

Beginning with the superset of 'Data processing systems or methods' (G06Q10) patents, we proceed to isolate genuine AIWM patents by analysing their full texts, sentence by sentence. The objective is to retrieve a representative collection of patents to gauge the intensity and extent to which the primary objectives identified in the literature are pursued. These objectives encompass the automation of direction, control, discipline, safety and health, in line with Coworker's identification of scopes of application. To achieve this, we employ a complex hierarchical text query that searches for the co-occurrence of a dictionary of target verbal predicates, direct objects and attributes within the same

[14] See: https://www.cooperativepatentclassification.org/sites/default/files/cpc/definition/G/definition-G06Q.pdf
[15] Coworker.org is a non-profit think-tank that deploys digital tools, data, and strategies to improve job safety and satisfaction, to help employees share information, form collectives, and advocate for change. See: https://home.coworker.org/worktech

sentence. Specifically, we investigate the presence of the three scopes identified in the EU-OSHA report:

- Increasing the efficiency and/or productivity of workers.
- Improving the decision-making process.
- Enhancing workers' health, safety and wellbeing.

However, it is important to note that these scopes are not always easily distinguishable and may exhibit some degree of overlap, particularly between the first and second scopes. The actual structure of the query is reported in Figure A 1. The rationale behind it is as follows: we aim to identify all potential combinations of text that indicate the presence of applications of technology designed to enhance worker efficiency and productivity, facilitate decision-making through automated practices, and enhance worker wellbeing and safety. The selected verbal predicates include terms such as 'assess', 'control', 'evaluate' and similar verbs. The primary objects of reference are human entities, encompassing 'labor', 'operator' and 'worker'. Moreover, we investigate various attributes, ranging from 'productivity', 'performance' and 'efficiency' to 'safety', 'emotion' and 'stress'. A total of 4,340 AIWM patents are singled out, and it is worth noting that the same patent may encompass multiple scopes within these defined parameters.

Once we have completed these steps, we can search for specific heuristics, like improving worker efficiency, enhancing decision-making, or promoting worker health and safety. We have developed a methodology to scan all the identified sentences and check for the co-occurrence of a specific verb, a direct object and an attribute that collectively convey the desired message within the same sentence. Figure A 1 displays the chosen words used in our query. In practice, we look for the simultaneous presence of a triplet of words (not necessarily adjacent), one from each set, within the same sentence. If at least one sentence contains at least one of these triplets (totalling 2,275 possibilities due to the combination of the three sets), we flag the associated patent as AIWM.

Figure A 1: Structure of the AIWM text query

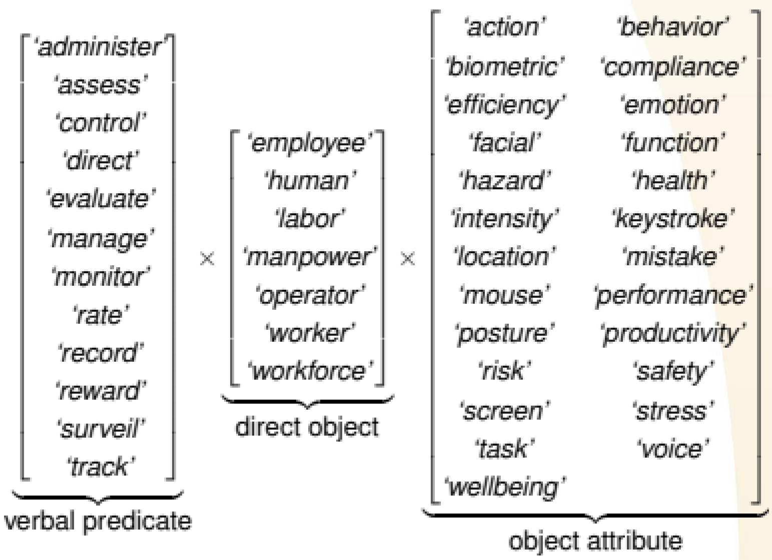

Source: Authors' elaboration

From AIWM patents to task and occupation exposure

Following the identification of the target patent pool, we link the textual content of these patents with tasks and human functions found in the O*NET dataset. This alignment forms the basis for constructing a quantitative measure of AIWM technological penetration, spanning from technology to labour markets. To achieve this, we aggregate tasks into occupations and occupations into sectors, utilising concordance tables provided by O*NET itself. This measure is established through pairwise text similarity between the corpus of patents and the corpus of task descriptions.

This methodological step draws significant inspiration from Montobbio et al. (2023). To align the technological content of AIWM patents with specific occupations, we utilise the O*NET database, which provides comprehensive descriptions of 19,265 distinct tasks, further grouped into 923 SOC2018 occupations using a weighting scheme that we will discuss later.

In the next steps, we aim to assess the textual similarity between the full texts of 4,340 AIWM patents and the 19,265 task descriptions. Each text is converted into a vector representing word frequencies. The number of components in each vector corresponds to the common dictionary of terms shared between these two large datasets. In simpler terms, all vectors exist within the same vector space, with dimensions equal to the number of distinct words in this shared dictionary. The similarity between each patent and task pair is then quantified by calculating the cosine of the angle between their respective vectors.

Instead of merely counting word occurrences in each text, we employ the tf-idf (term frequency–inverse document frequency) term-weighting scheme for computing relevant frequencies. This scheme is defined as follows.

Definition 1

Let D *be a collection of documents* d*, each composed of an ensemble of terms* t *from a dictionary* T*. The tf-idf measure of term* t *appearing in document* d *is defined as follows:*

$$\text{tf-idf}(t,d,D) := \text{tf}(t,d)\cdot\text{idf}(t,D), \forall d\in D, \forall t\in T$$

$$\text{tf}(t,d) := 1_d(t) = \begin{cases} 1 & \text{if} t\in d \\ 0 & \text{otherwise} \end{cases}, \forall d\in D, \forall t\in T$$

$$\text{idf}(t,D) := \log\left(\frac{|D|}{|\{d\in D: t\in d\}|}\right), \forall t\in T$$

The associated $|\mathrm{D}|\times|\mathrm{T}|$ *document term matrix* D^{D} *is an array of tf-idf measures for all documents* d *in the generic collection* D *and for all terms* t *in the relevant dictionary* T*. In other words,*

$$D^D_{d,t} = \text{tf-idf}(t,d,D), \forall d\in D, \forall t\in T$$

The tf-idf statistic serves as a measure of how significant a particular term is within a specific document when compared to other documents within a collection. The tf-idf value increases as a word appears more frequently within the document but is adjusted based on the number of documents in the entire corpus that mention the same word. This adjustment helps account for the fact that some words are more common overall.

Extending this concept to the corpus level, we create two document-term matrices, denoted as D^{AIWM} and D^{TASK}, in which the rows represent the tf-idf frequency vectors for each AIWM patent and each task description, respectively. Both matrices are based on the dictionary of terms derived from CPC definitions, which is the smaller of the two collections and consists of 6,868 terms. Consequently, D^{AIWM} has a dimension of 4340×6868 and D^{TASK} has a dimension of 19265×6868. In conclusion, we create the cosine similarity matrix denoted as S by calculating the cosine similarity scores between every pair of row vectors extracted from the document-term matrices D^{AIWM} and D^{TASK}. This calculation is done in accordance with the following definition.

Definition 2

Given two vectors $\mathrm{X},\mathrm{Y}\in\mathbb{R}^{+}$*, their cosine similarity is defined as the cosine of the angle between them, which is also equal to the inner product of the same vectors normalised to unit length, as follows:*

$$cos(X,Y) := \frac{X\cdot Y}{\|X\|\|Y\|} = \frac{\sum_{t=1}^{|T|} x_t\, y_t}{\sqrt{\sum_{t=1}^{|T|} x_t^2}\,\sqrt{\sum_{t=1}^{|T|} y_t^2}}$$

where $\mathrm{x_t}$ *and* $\mathrm{y_t}$ *denote the components of vectors* X *and* Y*, respectively, and* $\|\cdot\|$ *denotes the Euclidean norm.*

As the row vectors of document-term matrices consist of non-negative values, their cosine similarity falls within the bounds of the unit interval, meaning $cos(X,Y)\in[0,1]$. Furthermore, when term frequency is measured using tf-idf, the normalisation denominator in the previously mentioned equation becomes unnecessary, resulting in $cos(X,Y)\equiv X\cdot Y$. Therefore, when considering the document-term matrices

D^{AIWM} and D^{TASK} and extending the computation of cosine similarity to the matrix level, we arrive at the following cosine similarity matrix, which contains 83,610,100 values:

$$S = cos(D^{AIWM} \cdot D^{TASK}) \equiv D^{AIWM}(D^{TASK})'$$

Table A 1: Architecture of the cosine similarity matrix S

Occupation	11-1011.00			...	53-7121.00		
Patent \ Task	8823	8824	...	...	...	12809	12810
20010032103	cos(20010032103, 8823)	cos(20010032103, 8824)	...	...	...	cos(20010032103, 12809)	cos(20010032103,12810)
20010032110	cos(20010032110, 8823)	cos(20010032110, 8824)	...	...	...	cos(20010032110, 12809)	cos(20010032110, 12810)
...	...	...	...	...	...	...	...
20220207628	cos(20220207628, 8823)	cos(20220207628, 8824)	...	...	...	cos(20220207628, 12809)	cos(20220207628,12810)
20220210276	cos(20220210276, 8823)	cos(20220210276, 8824)	...	...	...	cos(20220210276, 12809)	cos(20220210276, 12810)

Source: Authors' elaboration

The architecture of the matrix is shown in Table A 1. Each row corresponds to one AIWM patent. Each column corresponds to one O*NET task. The top header groups multiple tasks in single occupations, in a many-to-one fashion. The aggregation of tasks into occupations is done in accordance with the O*NET classification of 'core' and 'supplemental' tasks. This distinction is in turn grounded on a blend of measures, namely 'importance', 'relevance' and 'frequency' of the task, relative to the underlying occupation, estimated by O*NET using survey data. In particular, we impute task similarity to occupations with the following weights:

$$\text{core: } \frac{2/3}{\text{\# tasks in the occupation}}$$

$$\text{supplemental: } \frac{1/3}{\text{\# tasks in the occupation}}$$

Robustness checks and external validation

As a robustness check, we assess the stability of the results when we systematically alter the text query. Through the repeated elimination of keywords, we aim to identify a patent pool that remains consistent regardless of specific terms used. To validate the set of AIWM patents and eliminate possible false positives, we apply a dependency parsing algorithm, which only retains (sentences of) patents that comply with robust syntactic structures. Additionally, we conduct two types of external validation exercises: (i) we compare our results with existing literature to ensure consistency; and (ii) we conduct two crucial interviews with OSH representatives from the firm heavily involved in developing these technologies.

Although the lists of predicates, direct objects and attributes presented in Figure A 1 appear suitable for capturing the desired heuristics, it is important to recognise that each individual word has the potential to introduce unintended biases into the final results. This is particularly relevant for words that carry some degree of ambiguity or are applicable to various technological domains. To mitigate any bias, we conducted a sensitivity analysis to assess how the set of AIWM patents and their cosine similarity to tasks changes when we remove one word at a time. Our findings confirm that none of the words used in the final query significantly skews the results.

Table A 2 provides information on the uniqueness of each word concerning AIWM patents (third column). Specifically, it indicates how many AIWM patents are uniquely associated with each word, meaning that if that word were excluded, the total count of AIWM patents would decrease by the same amount. Additionally, the table expresses this figure as a proportion of the final set of 4,340 AIWM patents (fourth column).

Table A 2: Sensitivity analysis of the AIWM text query

Type	Word	Unique	%	Type	Word	Unique	%
predicate	'administer'	17	0.39%	attribute	'compliance'	60	1.38%
predicate	'assess'	102	2.35%	attribute	'efficiency'	86	1.98%
predicate	'control'	450	10.37%	attribute	'emotion'	1	0.02%
predicate	'direct '	210	4.84%	attribute	'facial'	0	0.00%
predicate	'evaluate'	205	4.72%	attribute	'function'	359	8.27%
predicate	'manage'	337	7.76%	attribute	'hazard'	8	0.18%
predicate	'measure'	257	5.92%	attribute	'health'	96	2.21%
predicate	'monitor'	530	12.21%	attribute	'intensity'	4	0.09%
predicate	'rate'	213	4.91%	attribute	'keystroke'	2	0.05%
predicate	'record'	330	7.60%	attribute	'location'	617	14.22%
predicate	'reward'	66	1.52%	attribute	'mistake'	3	0.07%
predicate	'surveil'	0	0.00%	attribute	'mouse'	16	0.37%
predicate	'track'	395	9.10%	attribute	'performance'	662	15.25%
d. object	'employee'	1343	30.94%	attribute	'posture'	5	0.12%
d. object	'human'	604	13.92%	attribute	'productivity'	197	4.54%
d. object	'labor'	92	2.12%	attribute	'risk'	139	3.20%
d. object	'manpower'	32	0.74%	attribute	'safety'	71	1.64%
d. object	'operator'	1055	24.31%	attribute	'screen'	121	2.79%
d. object	'worker'	711	16.38%	attribute	'stress'	6	0.14%
d. object	'workforce'	42	0.97%	attribute	'task'	466	10.74%
attribute	'action'	333	7.67%	attribute	'voice'	64	1.47%
attribute	'behavior'	94	2.17%	attribute	'wellbeing'	1	0.02%
attribute	'biometric'	0	0.00%				

Source: Authors' elaboration

To maintain internal consistency in the sentences identified during the text query phase, we utilise a dependency parsing algorithm. Dependency parsing involves analysing the grammatical structure within a sentence to identify how words are interconnected and to determine the nature of the relationships between them. We draw inspiration from the work of Rughi et al. (2023), who conducted a similar exercise. Our requirement is for the sentences matched by the text query to conform to any of the following semantic patterns, for which we also provide an example (emphasis added):

1. predicate → attribute → object;

   'For example, … improved analysis of data from multiple data sources in order to **track** activities of and determine **productivity** of **employees**.' (US20220027426A1)

2. predicate ← attribute → object;

   '… the system has an opportunity to modify an internal operating state … based upon **recorded operator performance** …' (US20210339756A1)

3. predicate → object → attribute;

   'This information can then be used to **track** … **worker productivity**, or even pay and other harvest tracking.' (US20220083965A1)

4. object → attribute → predicate;

   'Identification of the operator is advantageous because it allows **operator performance** to be **monitored** and, if need be, allows operator behavior to be corrected.' (US20030164398A1)

5. object → predicate → attribute;

   'In addition, the … system enables **performance** data to be **tracked** for each worker and **task**.' (US20210350310A1)

Arrows → and ← denote relationships between words within the dependency tree of the underlying sentence. A sensitivity analysis similar to that carried out on the words of the text query reveals that:

- pattern 1. is unique to 1,795 AIWM patents (41.36%);
- pattern 2. is unique to 957 AIWM patents (22.05%);
- pattern 3. is unique to 287 AIWM patents (6.61%);
- pattern 4. is unique to 87 AIWM patents (2.0%); and
- pattern 5. is unique to 157 AIWM patents (3.62%).

References

Adams-Prassl, Jeremias (2022). “Regulating algorithms at work: Lessons for a ‘European approach to artificial intelligence’”. *European Labour Law Journal* 13(1), pp. 30–50. Available at: https://doi.org/10.1177/20319525211062558

Aloisi, Antonio, and Valerio De Stefano (2022). Your Boss Is an Algorithm: Artificial Intelligence, Platform Work and Labour. Bloomsbury Publishing.

Baiocco, Sara, Enrique Fernández-Macías, Uma Rani, and Annarosa Pesole (2022). The Algorithmic Management of Work and its Implications in Different Contexts. JRC Working Papers Series on Labour, Education and Technology No. 2022/02.

Ball, Kirstie (2021). Electronic Monitoring and Surveillance in the Workplace. European Commission Joint Research Centre.

Berg, Janine, and Pawel Gmyrek (2023). Automation hits the knowledge worker: ChatGPT and the future of work. UN Multi-Stakeholder Forum on Science, Technology and Innovation for the SDGs (STI Forum).

Bernhardt, Annette, Lisa Kresge, and Reem Suleiman (2021). Data and Algorithms at Work: The Case for Worker Technology Rights. UC Berkeley Labor Center. Available at: https://laborcenter.berkeley.edu/wp-content/uploads/2021/11/Data-and-Algorithms-at-Work.pdf

Bernhardt, Annette, Lisa Kresge, and Kung Feng (2023). Response to the White House Office of Science and Technology Policy Request for Information on Automated Worker Surveillance and Management. UC Berkeley Labor Center.

Cirillo, Valeria, Matteo Rinaldini, Jacopo Staccioli, and Maria Enrica Virgillito (2021). “Technology vs. workers: the case of Italy’s Industry 4.0 factories”. *Structural Change and Economic Dynamics* 56, pp. 166–183. Available at: https://doi.org/10.1016/j.strueco.2020.09.007

Delfanti, Alessandro (2021). “Machinic dispossession and augmented despotism: Digital work in an Amazon warehouse”. New Media & Society 23(1), pp. 39–55. Available at: https://doi.org/10.1177/1461444819891613

De Stefano, Valerio, and Simon Taes (2023). “Algorithmic management and collective bargaining”. Transfer: European Review of Labour and Research 29(1), pp. 21–36. Available at: https://doi.org/10.1177/10242589221141055

Dosi, Giovanni, Luigi Marengo, and Maria Enrica Virgillito (2021). “Hierarchies, Knowledge, and Power Inside Organizations”. Strategy Science 6(4), pp. 371–384. Available at: https://doi.org/10.1287/stsc.2021.0136

EU-OSHA – European Agency for Safety and Health at Work, Artificial intelligence for worker management: an overview, 2022a. Available at: https://osha.europa.eu/en/publications/artificial-intelligence-worker-management-overview

EU-OSHA – European Agency for Safety and Health at Work, Artificial intelligence for worker management: implications for occupational safety and health, 2022b, Available at: https://osha.europa.eu/en/publications/artificial-intelligence-worker-management-implications-occupational-safety-and-health

EU-OSHA – European Agency for Safety and Health at Work, Artificial intelligence for worker management: mapping definitions, uses and implications, 2022c. Available at: https://osha.europa.eu/en/publications/artificial-intelligence-worker-management-mapping-definitions-uses-and-implications

Gonfalonieri, Alexandre (2020). “What brain-computer interfaces could mean for the future of work”. Harvard Business Review, 6 October.

Mai, Jens-Erik (2016). “Big data privacy: The datafication of personal information”. The Information Society 32(3), pp. 192–199. Available at: https://doi.org/10.1080/01972243.2016.1153010

Massimo, Francesco (2020). “A Struggle for Bodies and Souls: Amazon Management and Union Strategies in France and Italy”. In: The Cost of Free Shipping: Amazon in the Global Economy, ed. by Jake Alimahomed-Wilson, Ellen Reese. Pluto Press.

Mast, Nina, and Lisa Kresge (2022). How Common is Employers’ Use of Workplace Management Technologies? A Review of Prevalence Studies. UC Berkeley Labor Center.

Montobbio, Fabio, Jacopo Staccioli, Maria Enrica Virgillito, and Marco Vivarelli (2022). “Robots and the origin of their labour-saving impact”. *Technological Forecasting and Social Change* 174, 121122. Available at: https://doi.org/10.1016/j.techfore.2021.121122

Montobbio, Fabio, Jacopo Staccioli, Maria Enrica Virgillito, and Marco Vivarelli (2023). “Labour-saving automation: a direct measure of occupational exposure”. *World Economy* [forthcoming].

Moore, Phoebe (2019). “OSH and the Future of Work: benefits and risks of artificial intelligence tools in workplaces”. EU-OSHA – European Agency for Safety and Health at Work. Available at: https://osha.europa.eu/en/publications/osh-and-future-work-benefits-and-risks-artificial-intelligence-tools-workplaces

Moro, Angelo, Matteo Rinaldini, Jacopo Staccioli, and Maria Enrica Virgillito (2019). “Control in the era of surveillance capitalism: an empirical investigation of Italian Industry 4.0 factories”. *Journal of Industrial and Business Economics* 46(3), pp. 347–360. Available at: https://doi.org/10.1007/s40812-019-00120-2

Negrón, Wilneida (2021). Little Tech Is Coming for Workers: A Framework for Reclaiming and Building Worker Power. Coworker.org.

Pesole Annarosa (2023). “Surveillance and monitoring of remote workers: implications for occupational safety and health”. EU-OSHA – European Agency for Safety and Health at Work. Available at: https://healthy-workplaces.osha.europa.eu/en/publications/surveillance-and-monitoring-remote-workers-implications-occupational-safety-and-health-0

Rughi, Tommaso, Jacopo Staccioli, and Maria Enrica Virgillito (2023). *Climate change and labour-saving technologies: the twin transition via patent texts.* LEM Working Papers series n. 11/2023.

Urzì Brancati, Cesira, and Maurizio Curtarelli (2021). Digital tools for worker management and psycho-social risks in the workplace: evidence from the ESENER survey. JRC Working Papers Series on Labour, Education and Technology No. 2021/12.

Urzì Brancati, Maria Cesira, Maurizio Curtarelli, Sara Riso, and Sara Baiocco (2022). How digital technology is reshaping the art of management. JRC Working Papers Series on Labour, Education and Technology, No. 2022/05.

Wood, Alex J. (2021). Algorithmic management consequences for work organisation and working conditions. JRC Working Papers Series on Labour, Education and Technology No. 2021/07.

Woodcock, Jamie (2022). “Artificial intelligence at work: The problem of managerial control from call canters to transport platforms”. Frontiers in Artificial Intelligence 5, 888817. Available at: https://doi.org/10.3389/frai.2022.888817

Authors: Jacopo Staccioli (Department of Economic Policy, Università Cattolica del Sacro Cuore) and Maria Enrica Virgillito (Institute of Economics, Scuola Superiore Sant'Anna).

Project Management: Emmanuelle Brun and Maurizio Curtarelli (EU-OSHA).