A Series on the EU AI Act

Pt. 4 – AI at Work

April 2025

Authored by
***Laura Lazaro Cabrera**, Counsel and Director of the Equity and Data Programme, CDT Europe*
***Magdalena Maier**, Legal and Advocacy Officer, Equity and Data Programme, CDT Europe*

In the past years, the use of algorithmic management and decision-making systems in the workplace has become more and more widespread: a recent OECD survey found that over 70% of consulted managers reported that their firms used at least one automated tool to instruct, monitor or evaluate employees. This increase in use is understandably being met with apprehension. A survey conducted this year by the European Commission underscores workers' overwhelming support for rules regulating the use of AI in the workplace, endorsing the European Trade Union Confederation's previous calls for a Directive on algorithmic systems in the workplace that would specifically tackle some of the emerging challenges.

The EU's AI Act, the first cross-cutting landmark regulation of AI, recognises the risks involved in the deployment of AI systems in the workplace and it creates specific obligations aimed at ensuring the protection of workers through prohibitions and increased safeguards, with varying levels of success.

Building on the previous explainers in this series, this brief zooms in on the specific aspects of the AI Act that are most relevant in the context of employment and the rights of workers in light of existing EU legislation on the protection of workers.

This explainer will focus on the obligations of employers using AI systems in the workplace. Under the AI Act taxonomy, employers using AI will qualify as deployers of an AI system, regardless of whether an AI system is developed in-house – in which case they could be considered to be both providers and deployers – or acquired for use in the workplace.

Prohibited AI systems: red lines in the employment context

In line with its risk-based approach, the AI Act prohibits several AI practices which it considers to pose an unacceptable risk – several of which directly or indirectly are relevant to the workplace. While only a prohibition on the use of emotion recognition systems in the workplace explicitly relates to the employment context, several other prohibited AI systems have the potential to adversely impact the rights of workers, such as biometric categorisation systems or social scoring systems. We explore the prohibitions with the most salient impacts on the workforce below, in order of strength.

Biometric categorisation - entirely prohibited

The Act prohibits AI systems which categorise individuals based on their biometric data to deduce or infer a series of attributes, including race, political opinions, and trade union membership among others (Article 5(1)(g)). This prohibition captures an employer relying on biometric categorisation to find out whether an individual belongs to a specific trade union, which could lead to negative consequences for that individual worker. This prohibition could similarly be relevant in the context of recruitment, for example if a job advertisement is only shown to certain groups of people based on their prior categorisation.

Emotion recognition - (Mostly) prohibited in employment settings

Acknowledging the well-established unreliability of emotion recognition systems (Recital 44), the AI Act prohibits the placing in the market and use of AI systems that infer emotions from individuals in the workplace, except when such systems are put in place for medical or safety reasons (Article 5(1)(f)). Emotion recognition under the Act is defined not in terms of an AI system's capability, but in terms of its purpose, namely "identifying or inferring emotions or intentions of natural persons on the basis of their biometric data". The Act excludes from the definition systems to recognise physical states, such as pain or fatigue (Recital 18), which are otherwise permitted.

The guidelines on prohibited AI practices issued by the EU AI Office provide key clarifications on the scope of the prohibition. First, the guidelines apply a broad interpretation of "workplace", clarifying that the prohibition extends to the recruitment process – in other words, job applicants or candidates are protected even in the absence of a formal employment or contractual relationship. Second, the guidelines clarify that the exception for medical and safety reasons should be interpreted narrowly, with any proposed interventions being required to be (i) responsive to an explicit need, (ii) limited to what is "strictly necessary", including limits in time, personal application and scale, and (iii) accompanied by sufficient safeguards. Consequently, the guidelines specify that the "medical reasons" exception cannot be relied upon to legitimise the detection of general aspects of wellbeing, including monitoring of stress levels. Likewise, "safety reasons" pertain only to the protection of life and health, and cannot be relied upon to legitimise the use of emotion recognition for the purposes of protecting property interests, for example to protect against theft or fraud.

Despite the welcome clarifications above, the guidelines introduce carve-outs not foreseen in the text of the prohibitions itself. Notably, they exclude systems deployed for personal training purposes as long as the results are not shared with human-resources responsible persons and cannot impact the work relationship of the person trained or their professional progression. This carve-out enables employers to require workers to undergo emotion recognition for training purposes - even if the results are not shared, a third-party company contracted to provide such training could inform the employer whether such training was undertaken or not. Moreover, the guidelines state that crowd-control measures in public spaces continue to be allowed even if this means that employees present in the area will be subject to emotion recognition, given that this is not the primary aim of the measure. Consequently, employees working for example at a sports stadium could still be lawfully subject to emotion recognition according to the guidelines.

Social scoring - prohibited on a case-by-case basis

Furthermore, the AI Act prohibits systems used for social scoring of individuals or groups based on their social behaviour or known or inferred characteristics whenever the score leads to detrimental treatment in an unrelated context or to detrimental treatment disproportionate to the social behaviour or its gravity (Article 5(1)(c)). In the workplace context, the latter is likely to be more relevant, and could include situations where a worker is fired or demoted based on their behaviour and inferred personality traits – such as perceived introversion or aloofness – such that treatment is unjustified or disproportionate to the social behaviour itself or its gravity. However, whether or not the practical consequence of a poor social scoring results in disproportionate treatment will likely ultimately turn on the facts of the specific case at hand. In this regard, it is crucial to note that the Act itself notes that the social scoring prohibition does not apply to lawful evaluation practices carried out for a specific purpose (Recital 31), and the guidelines on prohibited practices specifically cite specific employee evaluations as an example of lawful evaluation practices, noting that "they are not per se prohibited, if lawful and undertaken in line with the AI Act and other applicable Union law and national law". The guidelines therefore signal that the use of social scoring in worker evaluations is not de facto prohibited, while cautioning that it could fall foul of the AI Act if all elements of the prohibition were met.

Real-time biometric identification - permitted

Finally, the AI Act prohibits real-time remote biometric identification specifically in the context of law enforcement (Article 5(1)(h)), implicitly acquiescing to the lawfulness of its use whenever used for purposes other than law enforcement. Such systems can therefore potentially be lawfully introduced and used by the employer to surveil workers under the AI Act, even as they might be subject to restrictions under the General Data Protection Regulation or other laws.

Limited protections from high-risk systems

The bulk of the AI Act is dedicated to regulating the development and deployment of high-risk AI systems, which are overall permitted but subject to safeguards, ranging from general notice requirements to the availability of effective remedies.

An AI system can be considered high-risk under the Act if it is listed in Annex III of the Act. This includes systems deployed in employment and self-employment, in particular i) recruitment and selection, ii) promotions and termination, iii) allocation of tasks and monitoring and iv) evaluation of performance (Annex III 4(a)).

As we have commented numerous times, one of the key shortcomings of the Act is that it allows the possibility for an AI system deployed in any of the settings described in Annex III – including those set out above – to escape the high-risk classification if it is considered that a given system does not pose a significant risk of harm to the health, safety or fundamental rights of individuals (Article 6(3)). If a system is not recognised as being high-risk by a provider, most of the AI Act obligations are inapplicable – including those pertaining to deployers. Nevertheless, providers deeming an AI system not to be high-risk despite being covered by Annex III are asked to document this assessment (Article 6(4)), and register their system in a publicly available database (Article 49(2)). The AI Act further requires deployers who are public authorities not to use a high-risk AI system if it has not been listed by a provider in the publicly available database, creating an additional safeguard for their employees (Article 26(8)), but no similar restriction operates for private sector employees.

The high-risk classification is essential for key fundamental rights protections to kick in. High-risk systems are subject to risk management obligations, which include the identification of risks that the high-risk AI system can pose to health, safety or fundamental rights, transparency obligations towards deployers, and guarantees relative to human oversight, among others.

Deployers of a high-risk AI system – which includes employers – specifically have several key obligations enabling the transparency and accountability of the use of AI systems in the workplace. These obligations vary based on the identity of the deployer.

Obligations applying to all deployers

The AI Act imposes general obligations on deployers, including ensuring some level of human oversight and monitoring the functioning of an AI system.

Where the workplace is concerned, the AI Act creates a concrete notice obligation for deployers, requiring deployers of high-risk AI systems to inform workers' representatives and affected workers that they will be subject to an AI system prior to putting such a system in place (Article 26(7)). The recitals leave the door open to go beyond mere notice requirements, noting that the Act is without prejudice to worker consultation procedures laid down in EU law (Recital 92) – however existing laws cover consultation procedures in a patchwork manner. The Workers' Safety and Health Directive requires consultation with workers and/or their representatives on the planning and introduction of new technologies, specifically regarding the consequences of the choice of equipment, the working conditions and the working environment for the safety and health of workers (Article 6(3)(c)). The Directive on informing and consulting employees obliges employers beyond a given size to consult with their employees on decisions likely to lead to substantial changes in work organisation, while leaving the regulation of the practical arrangements to the Member States (Article 4(2)(c)). Consequently, this Directive has the potential to cover a wider scope of AI systems with implications for workers' rights, besides their safety and health. Nevertheless, it is unclear whether the introduction of AI would fall within Member States' definition of "substantial changes."

The consultation obligation set out in Directive 2002/14/EC has been interpreted by the recently adopted Platform Work Directive to include "decisions likely to lead to the introduction of or to substantial changes in the use of automated monitoring systems or automated decision-making systems" (Article 13(2)). This Directive also regulates in detail the information digital labour platforms need to provide to platform workers, their representatives and national competent authorities in the context of automated monitoring and decision-making systems (Article 9). It is, however, important to keep in mind that this Directive only applies to work organised through a digital labour platform (Article 2(1)(a) and (b)). This includes work performed completely online, including professional tasks such as software development or translation services, or in a hybrid manner combining online communication with a real-world activity, for instance the provisions of transportation services or food delivery (see Recital 5). It therefore remains to be seen to what extent the obligation to consult under Directive 2002/14/EC also applies to regular workspaces.

From a rights perspective, consultations are only the starting point – how they are conducted, and the extent to which the results are taken on board are crucial to ensure their effectiveness. The AI Act leaves the possibility for more favourable legislation for workers in the Union or Member States open (Article 2(11)). Consequently, for instance, whether workers or their representatives have a veto over the introduction of AI systems depends on the national law and collective agreements in place.

Obligations applying to deployers who are public authorities or perform public services

The AI Act creates additional obligations for deployers who are public authorities, which are held to a higher standard. As already explored above, public authorities cannot deploy a high-risk AI system that has not been previously identified and registered as such by a provider in a public database. Further, the Act requires public authorities to conduct a fundamental rights impact assessment (FRIA) prior to the deployment of an AI system identified as high-risk in Annex III (Article 27) and the registration of a high-risk AI system being used in a publicly available database (Article 26(8)). While these obligations are crucial in ensuring the transparency and accountability of use of an AI system in the workplace, there are important nuances to be taken into account.

The obligation to conduct a FRIA applies not only to entities governed by public law, but also – crucially – to private entities performing public services, which the AI Act considers to cover entities providing services "linked to tasks in the public interest", such as in the areas of education, healthcare, social services, housing, and the administration of justice (Recital 96). The list provided is non-exhaustive, opening up the possibility for entities performing other functions to be covered. FRIAs are a unique feature and perhaps the most positive aspect under the AI Act. Unfortunately however, this obligation only applies in the narrow circumstances identified above, meaning that the majority of private employers are not required to assess the impact of the AI system's use on the fundamental rights of their employees before deployment. Once conducted, there is no obligation on the employer to disclose the full results of the FRIA beyond notifying the national regulator of the outcome, limiting the potential for employee awareness and oversight.

Beyond conducting a FRIA, the AI Act requires public sector deployers or any entity acting on their behalf to register any high-risk AI systems used in a public database, providing basic information on the AI system in an accessible manner (Article 71), and specifically including a summary of the fundamental rights impact assessment and data protection impact assessment (Annex VIII Section C). On this basis, workers could expect to see a brief summary of any anticipated fundamental rights impacts, as well as any mitigations undertaken by their employer.

Remedies, enforcement and governance

As explained in a previous blog post, the AI Act contains only a limited number of remedies, which are solely available for individuals having been subjected to a high-risk AI system within the meaning of Annex III. These remedies consist of the right to an explanation for a decision taken based on the output of a high-risk AI system, as well as the right to lodge a complaint.

The AI Act gives individuals subject to a decision based on a high-risk system's output the right to a clear and meaningful explanation by the deployer of the system (Article 86), building on the right not to be subjected to automated decision-making (ADM) with legal or similar effects on individuals, laid down in the General Data Protection Regulation (GDPR). The GDPR further requires the data controller to inform individuals about the existence of automated decision-making, the logic involved as well as the significance and consequences of such processing (Articles 13(2)(f) and 14(2)(g)). Where GDPR creates a base layer of protection shielding individuals from the serious consequences of automation, the AI Act introduces an additional dimension of protection by entitling individuals to information about consequential decisions taken not solely through automated means, but nonetheless relying on its support.

The right to a clear and meaningful explanation can be a useful tool for employees to open up the "black box" of an algorithmic management or decision-making system and understand its logic, potentially enabling them to assess whether they have been adversely affected. However, the Act is not clear whether the explanation is to be provided proactively or whether individuals are entitled to receive it only upon request. In the latter case, the burden would be on employees to remain alert to any decisions likely taken with the support of AI systems. Further, as most employers will probably struggle to fully comprehend the logic of the AI system themselves, such explanations may be inaccurate or incomplete and will therefore not always contribute to a better understanding of the situation. Lastly, the explanation - if meaningfully given - is no guarantee of corrective action, which will have to be sought outside of the scope of the AI Act.

The AI Act creates the right for any individual to lodge a complaint before a national market surveillance authority if they consider any part of the AI Act has been infringed, regardless of whether they have been personally affected or not (Article 85).

For example, an employee could bring a complaint if:

- they did not receive an explanation for a decision taken based on the output of a performance-monitoring AI system at work;
- their public sector employer deployed a high-risk AI system at the workplace without disclosing it in the public database of AI systems; or
- their private sector employer failed to give prior notice to the workforce about a high-risk AI system being rolled out at work.

As we have previously analysed, the right to lodge a complaint is limited as it does not include an obligation for a national authority to investigate or to respond. Nevertheless, it is an additional regulatory avenue for individuals suspecting foul play and any violation of the AI Act.

The AI Act creates several oversight mechanisms to invite sector-specific expertise in the enforcement of the AI Act. Notably, the AI Act provides for the designation of fundamental rights authorities at national level who may request and access documentation created in observance of the obligations of the AI Act in accessible language and format to exercise their mandate (Article 77(1)). In some Member States, those authorities include institutions active in the context of workers' rights and labour law, such as labour inspectorates or occupational health and safety institutions. These authorities can therefore ask for the necessary information on the deployed AI system to facilitate the exercise of their mandate and protect the rights of workers.

Finally, the AI Act establishes an Advisory Forum to provide technical expertise and advice with a balanced membership from industry, start-ups, SMEs, civil society and academia. While there is no explicit involvement of social partners on it, the Forum could provide an important platform for stakeholders to specifically bring in the perspectives of workers and their rights.

Conclusion

In conclusion, while the AI Act's minimum harmonisation approach in the context of employment is a positive step, allowing more favourable laws to apply, the regulation itself has only limited potential to protect workers' rights – with its main contributions being the restriction of the use of emotion recognition in the workplace, creation of notice obligations and explanation mechanisms. In particular, the obligations of employers deploying high-risk systems come with significant loopholes and flaws. Likewise, workers and their representatives have limited remedies available in the case of AI-induced harm. Potential secondary legislation could strengthen workers' rights to be meaningfully consulted before the introduction of algorithmic management and decision-making tools. It should furthermore require all employers to consider the fundamental rights impact of those systems and ensure their transparency and explainability to workers and their representatives.

As the AI Act is gradually implemented, important aspects to monitor are the use of notice and – where applicable under existing EU or national law – consultation mechanisms at the worker level, as well as the interpretation and operationalisation of the right to obtain an explanation. Another crucial area of inquiry will be the extent to which private entities can be considered to be providing public services on a case-by-case basis. It is therefore vital that CSOs and workers' rights organisations are meaningfully engaged in the AI Act's implementation and enforcement processes.

Find more from the CDT Europe team on the EU AI Act at *cdt.org/eu*.

*The **Center for Democracy & Technology (CDT)** is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1994, CDT has been a trusted advocate for digital rights since the earliest days of the internet. The organization is headquartered in Washington, D.C. and has a Europe Office in Brussels, Belgium.*