Alessandra Ingrao

AI at Work: Reframing Data Protection through the Lens of Labor Law

Contents: **1.** Introduction. **2.** Against AI exceptionalism: a methodological warning for data protection and labor law. **3.** Two structural flaws in the approach of privacy law in the age of AI. **4.** Between continuity and disruption: the AI act and the risk of normative exceptionalism. **5.** The erosion of purpose limitation in the age of adaptive AI: from determinism to opacity. **6.** Challenging the principle of data minimization in ai-driven workplaces. **7.** Automated decision-making in the workplace: limits of individual rights and the need for collective oversight.

1. *Introduction*

Artificial Intelligence (AI) has become an integral component of contemporary productive and organizational dynamics, profoundly reshaping the modalities of work performance, the exercise of managerial authority, and workplace control techniques[1]. In this evolving context, the protection of workers' personal data and, more broadly, the safeguarding of their private sphere – emerges with renewed centrality, raising complex normative and systemic questions. This essay seeks to elucidate the key tensions between data protection law and the deployment of AI in the employment context, with the aim of critically assessing the adequacy of existing legal instruments and identifying potential regulatory trajectories.

[1] Alaimo, *Il Regolamento sull'Intelligenza Artificiale. Un treno al traguardo con alcuni vagoni rimasti fermi*, in *Federalismi*, 2024, p. 231 ff.; L. Zoppoli, *Il diritto del lavoro dopo l'avvento dell'intelligenza artificiale: aggiornamento o stravolgimento? Qualche (utile) appunto*, in *DLM*, 2024, 3, p. 1 ff.; Ciucciovino, *Intelligenza artificiale e diritto del lavoro: problemi e prospettive*, in *DRI*, 2024, 3, p. 586 ff.; Nuzzo, *Vecchi e nuovi limiti al monitoraggio dei lavoratori al tempo dell'IA*, in *RGL*, 2024, 4, p. 555 ff.

To engage meaningfully with the intersection of AI and the protection of workers' rights, it is first necessary to clarify what is meant by "artificial intelligence" today. In its current usage, the term refers to a heterogeneous set of computational tools primarily grounded in machine learning techniques, including deterministic, non-deterministic, and generative models. Far from the science fiction image of sentient robots, contemporary AI systems are algorithmic constructs designed to process vast quantities of data and generate inferences, predictions, or new forms of content. They do not represent a paradigmatic break with the past, but rather a successful recombination of existing technologies, whose operational effectiveness has been enhanced by exponential advances in computational power and data availability.

Nevertheless, the label "artificial intelligence" has acquired a powerful symbolic role in public and regulatory discourse, functioning as an organizing metaphor that attracts attention, resources, and normative legitimacy. As a result, a technology that is neither truly "intelligent" nor wholly "artificial" has acquired disproportionate symbolic prominence. These systems do not "learn" in any human sense; rather, they detect patterns and correlations in data selected, annotated, and structured by human agents[2]. Behind the façade of automation lies an extensive network of human labor – often – invisible that sustains AI's operational viability. In the world of work, acknowledging this reality is essential: the adoption of AI in personnel selection, performance evaluation, shift allocation, or predictive surveillance continues a longstanding trajectory of technological rationalization of employer power – one already familiar to labor law and requiring renewed critical engagement.

[2] The legal definition of an artificial intelligence system, set out in Article 3, par. 1(1) of the AI Act, confirms this premise. It describes an AI system as an automated system designed to operate with varying levels of autonomy and capable of producing outputs such as predictions, content, recommendations, or decisions that may influence physical or virtual environments. Crucially, Recital 12 emphasizes the system's inferential capacity – that is, its ability to derive models or algorithms from data inputs and generate outputs that exceed basic data processing, enabling learning, reasoning, or modelling (see Recital 12). This is precisely what distinguishes AI from traditional rule-based software, which follows predefined instructions without the capacity for autonomous decision-making or adaptation over time. Regarding the complexity of defining an AI system, see also the European Commission's report summarising the responses of stakeholders to the public consultation. Commission Guidelines on the definition of an artificial intelligence system, 6 February 2025, C (2025) 924 final.

Secondly, it is crucial to understand the specific privacy challenges posed by AI[3]. These concern both the input phase (data collection and selection) and the output phase (generation of inferences, analysis, classifications, and decisions). Practices such as non-consensual data scraping or the large-scale harvesting of nominally legitimate data often evade the protections established by current legal frameworks. On the output side, the use of algorithms to derive information not explicitly provided by workers, to evaluate performance, or to predict future behaviors introduces unprecedented scenarios of profiling and control. These processes risk undermining the dignity of the worker, circumventing privacy safeguards, and exacerbating manipulation and surveillance risks.

AI also tends to replicate and reinforce preexisting systemic biases, contributing to the depersonalization of decision-making and eroding worker autonomy. The technical opacity of AI complicates transparency and accountability[4], making it difficult for affected individuals to understand or contest the decisions that affect them. Finally, the strategic economic value of AI technologies encourages deregulatory development paths in which the protection of fundamental rights may be subordinated to the imperatives of innovation and competitiveness.

In sum, AI does not represent a radical rupture but rather an acceleration of longstanding dynamics. Yet, precisely because of its capacity to intensify preexisting issues, it starkly exposes the structural gaps and ambiguities of current privacy regimes. This essay thus offers a critical examination of the principal legal challenges, and – drawing also on the normative legacy of labor law – proposes regulatory strategies capable of safeguarding human dignity and autonomy in an increasingly "datafied" workplace.

[3] On these aspects, see EDPB Opinion 28/2024, adopted on 17 December 2024 pursuant to Article 64(2) GDPR, which addresses critical issues such as the anonymisation of AI models, the use of legitimate interest as a legal basis for the development and deployment of such models, and the consequences of processing unlawfully obtained data during training on the model's subsequent lawfulness. The EDPB emphasises the need for a case-by-case assessment of whether an AI model can be considered anonymous, reiterates the requirement to apply the three-part legitimate interest test (necessity, proportionality, balancing of interests), and affirms that the illegality of training data may compromise the lawfulness of the model itself, unless proper anonymisation has been achieved.

[4] Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information*, Harvard University Press, 2015.

2. *Against AI exceptionalism: a methodological warning for data protection and labor law*

When addressing issues related to privacy and artificial intelligence, it is essential to resist what has aptly been described as "AI exceptionalism"[5] – the growing tendency in legal and policy discourse to frame AI as a radically novel, pervasive, and unpredictable technology that requires a separate and autonomous regulatory framework. Such a perspective risks distorting both the interpretive and normative landscape.

In reality, the concerns raised by AI – privacy violations, surveillance, lack of transparency, and discrimination – are not unprecedented. Rather, they are more extreme, complex, and opaque manifestations of longstanding issues that have already been addressed, albeit imperfectly, by existing regulatory instruments, particularly in labor law. AI has not created these problems; it has merely intensified them, made them more urgent, and harder to ignore.

Rising societal and institutional anxiety about AI has led European policymakers to propose new legislation, such as the AI Act, which establishes dedicated regulatory agencies and outlines governance frameworks that are structurally distinct from existing data protection regimes. Yet the core question is not whether a new law is needed, but whether lawmakers possess a sufficiently comprehensive and accurate understanding of the problems AI poses, and whether they are capable of identifying their true nature. A general or overly procedural – regulation such as the current European approach – risks overlooking the very substantive dimensions that privacy and labor law already seek to govern, despite their limitations.

Artificial intelligence must be understood as part of the broader historical trajectory of the digital transformation of labor relations – an evolution marked by the exponential growth in data collection, processing, and profiling. This trajectory has already prompted regulatory responses through instruments such as the GDPR, national labor statutes, anti-discrimination laws, occupational health and safety regulations, and the recent introduction of algorithmic transparency obligations under EU and domestic law.

It would thus be misguided to assume that privacy and labor law are already fully equipped to handle the challenges of AI, or that all that is needed

[5] Solove, *Artificial Intelligence and Privacy*, in *FLR*, 2025, vol. 77, p. 1 ff.

is an additional layer of protections. That would be akin to building a new floor on an already unstable foundation. At the same time, we must not start from scratch. What is required is a structural reconsideration of existing regulatory paradigms – a critical re-evaluation that acknowledges real discontinuities without neglecting deep continuities.

AI is not a parallel universe, but rather the continuation and intensification of processes that the law has long engaged with and, in part, already regulates.

In the field of labor, this means that the challenges posed by artificial intelligence must be addressed in light of the protections already in place. These protections are not necessarily obsolete, but they require updating, integration, and realignment. In this context, the principle of complementarity set forth by the AI Act plays a crucial role, outlining a regulatory framework that is minimal and non-exhaustive: "minimal" because it does not preclude the adoption of more favorable measures for workers at the national level, including through collective bargaining (Art. 2, §11); and "complementary" because it is not intended to undermine existing EU or national legal frameworks, but rather to operate functionally, facilitating and supporting existing rights and remedies (Recital 9)[6]. The very structure of the AI Act thus rejects an exceptionalist approach and reinforces the need for clear, coherent, and harmonized regulation. The real risk does not lie in the absence of new norms, but in losing direction – chasing the illusion of normative exceptionalism rather than strengthening, evolving, and rendering fully effective the legal framework of labor law in the digital age.

In light of these reflections, the following section identifies and analyses the most pressing challenges currently emerging at the intersection of AI, data protection, and labor law.

3. *Two structural flaws in the approach of privacy law in the age of AI*

The first major flaw in contemporary data protection architecture lies in its continued reliance on a model of individual informational control.

[6] This is expressly confirmed in the provisions concerning the deployer. In particular, it is clarified that the obligation to use the system in accordance with the provider's instructions must not compromise compliance with obligations established by other legal sources (Art. 26, §3, AI Act).

Since its inception, privacy law has largely been built upon the notion that empowering individuals through access to information, consent mechanisms, and post hoc rights – such as access, rectification, and objection – would suffice to safeguard personal autonomy in the digital age.

This logic has increasingly informed labor regulation as well. In recent decades, traditional labor law protections – such as the prohibitions under Article 4 of the Italian Workers' Statute against employer surveillance – have given way to more transparency-based frameworks. Notably, Article 1-bis decree n. 152/1997 embodies a shift from categorical prohibitions to a system premised on prior individual information. Under this model, it is assumed that if a worker is adequately informed about the source of the data collected and the logic of the algorithmic systems used for monitoring or decision-making, they will be better equipped not only to align their conduct with the employer's expectations, but also to exercise their rights more effectively, act autonomously, and contribute responsibly to organizational life.

Yet surprisingly little critical attention has been paid to this model, which continues to rest on the increasingly tenuous assumption that fully informed individuals can meaningfully navigate the complexities of data processing[7]. In practice, workers rarely read privacy notices, and when they do, they are often left with a sense of opacity and powerlessness. Even when privacy statements are read and understood, such awareness proves largely ineffective, as it does not translate into any actual capacity to influence the power structure overseeing data use. Thus, regulatory provisions grounded in the ideal of the "empowered data subject" often reveal their conceptual fragility, exposing a normative framework that remains markedly individualistic in orientation.

This is precisely where labor law – rooted in solidaristic and collective logics – can offer a corrective[8]. It reminds us that power asymmetries in the workplace are not resolved through information alone, but require mechanisms of participation and representation capable of articulating collective interests. Privacy governance in the workplace, therefore, cannot rely exclusively on individual empowerment; it must also incorporate institutionalized forms of worker voice and negotiation.

[7] See Hartzog, Richards, *Privacy's Constitutional Moment and the Limits of Data Protection*, in *Bost. Coll. LR*, 2020, 61, p. 1687 ff.

[8] Corti, *La partecipazione dei lavoratori: avanti piano, quasi indietro*, in Id (a cura di), *Il pilastro europeo dei diritti sociali e il rilancio della politica sociale dell'UE*, Vita e pensiero, 2021, p. 163 ff.

A second, and perhaps deeper, structural flaw in the current regulatory framework lies in its accountability model[9] – an architecture that appears increasingly misaligned with the operational logic of artificial intelligence. Similar to what has occurred in the field of occupational health and safety – where physical or psychological risks arise from work organization – the GDPR views "informational-technological risk" as a direct consequence of adopting digital tools capable of collecting, processing, and utilizing personal data. The GDPR marks a significant evolution beyond the consent-based paradigm, shifting the focus toward the proactive duties of data controllers. These duties – ranging from conducting data protection impact assessments and maintaining records of processing activities to ensuring data minimization and embedding data protection by design and by default – aim to bind organizational conduct to the effective protection of fundamental rights, including those of workers.

However, what is often overlooked in both academic and policy discussions is whether this accountability-based framework remains viable in the face of AI's expansive data demands. Unlike traditional monitoring systems (such as CCTV), AI requires access to substantially larger datasets to function effectively. More importantly, algorithmic decisions are no longer based solely on the data of individual subjects, but on inferential patterns drawn from the aggregation of data across millions of individuals[10]. In this context, the principle of data minimization is not merely difficult to apply – it risks becoming conceptually irrelevant. The issue is not one of non-compliance, but rather of a structural incompatibility between the principle's intent and the technological requirements of AI systems.

Moreover, it is important to highlight that risk analysis and mitigation strategies in the architecture of GDPR remain a unilateral obligation of the

[9] This principle requires the employer, as the data controller, to assess risks in advance and adopt appropriate measures to prevent or mitigate negative consequences. Regardless of the operational autonomy of processing systems, deployers are required to demonstrate the adoption, effective implementation, and continuous monitoring of the risk prevention model. These measures must be documented in the data protection impact assessment (Article 35 GDPR), describing the prevention strategies adopted in accordance with the processing principles set out in Article 5 GDPR. This article represents the only mandatory requirement for the employer to examine technological risks and the measures undertaken to mitigate them.

[10] S. Brown, *Machine Learning, Explained*, in *MIT Sloan School of Management*, 21 april 2021, https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained (last access: 3 may 2025).

data controller, carried out without mandatory participation from trade union representatives of workers. Article 35(9) of the GDPR considers obtaining the opinion of representatives of the affected categories as merely optional.

The AI Act fails to correct the unilateral and individualistic approach that characterises the regulation of artificial intelligence. The Fundamental Rights Impact Assessment (FRIA), as set out in Article 27, is mandatory only for public bodies and private entities providing services of general interest, such as schools, hospitals, or banks. However, it does not apply to deployers who use AI systems in the fields of "employment, worker management, and access to self-employment". During the final approval phase of the Regulation, the provision that would have imposed such an obligation – alongside essential safeguards such as human oversight and consultation – was removed. These elements had formed the protective core of the original legislative proposal (see former Article 29 bis).

Conversely, Article 8 of the Digital Platforms Directive recognizes trade union participation as an added value, imposing an obligation to consult workers and their representatives during risk assessments[11]. While employers are not bound to follow these opinions, the principle of accountability requires them to justify any deviations from the received feedback, ensuring such opinions are documented in the data protection impact assessment.

Another relevant aspect concerns the confidentiality of documents held by employers, an issue frequently encountered both in the privacy impact assessment and in delivering the Risk Assessment Document (DVR) to worker safety representatives. In this regard, Article 8(2) of the Digital Platforms Directive explicitly mandates the delivery of the impact assessment to worker representatives – a requirement absent in both the GDPR and the AI Act.

In this light, AI exposes a latent tension already present in privacy law – between the *ex ante* logic of restraint and a digital infrastructure that is, by design, driven by the continuous expansion of data collection and processing. This contradiction must be squarely confronted if data protection is to remain meaningful in the algorithmic workplace.

[11] Delfino, *Lavoro mediante piattaforme digitali, dialogo sociale europeo e partecipazione sindacale*, in *Federalismi.it*, 2023, 25, p. 171 ff.

4. *Between continuity and disruption: the AI Act and the risk of normative exceptionalism*

Despite the explicit acknowledgment of the AI Act's complementary and minimal character, which affirms the continued validity of pre-existing national and European legislation and encourages their integration (Recital 9 and Article 2(11)), the Regulation still risks, in certain respects, reinforcing the very form of regulatory exceptionalism that ought to be avoided. The Act often treats artificial intelligence as a technology requiring an autonomous and distinct legal framework, rather than as a phenomenon to be governed within existing legal paradigms – chief among them, data protection law. Such an approach may lead to excessive fragmentation, regulatory duplication, and conceptual misalignment, particularly in the domain of labor, where strong safeguards are already in place[12].

This tendency is especially problematic when viewed through the lens of two foundational data protection principles: *purpose limitation* and *data minimization*. The proper functioning of AI systems – especially those relying on machine learning techniques – often presupposes the ingestion and processing of large, heterogeneous datasets. In many cases, the utility and accuracy of such systems increase with the volume and diversity of data they can access. Yet this requirement stands in tension with legal obligations to collect only data that is necessary and relevant for specific, clearly defined purposes. The expansionist logic of AI thus places strain on these core principles, calling into question whether the current legal architecture is structurally equipped to manage such a conflict.

High-risk AI systems, as defined under the AI Act, are subject to a range of ex ante compliance obligations imposed primarily on producers[13]. These include conformity assessments, CE markings, technical documentation, and

[12] *Funditus* M.T. Carinci, Ingrao, *L'impatto dell'AI Act sul diritto del lavoro*, in *DLRI*, 2024, 184, p. 451 ff.

[13] Article 5 AI Act sets out a list of prohibited practices, including emotion recognition in the workplace, untargeted scraping of biometric data, and subliminal manipulation (AI Act, Art. 5, par. 1, a-g), see M.T. Carinci, Ingrao, *cit.*, p. 463. The European Commission's Guidelines on Prohibited AI Practices, published on 4 February 2025, provide detailed interpretations of each prohibited practice and clarify their scope of application. They confirm, among other things, that "deployers" (i.e., employers) are also subject to the prohibitions set out in Article 5. Although non-binding, these guidelines offer valuable interpretative guidance and serve as best practices supporting regulatory enforcement.

risk management protocols. These obligations, modelled on product safety and liability regimes, are intended to ensure that AI systems entering the European market meet defined technical and ethical standards[14]. However, in practice, many of these safeguards rely on internal compliance mechanisms – especially self-assessment by providers[15] – rather than oversight by independent third-party bodies. Notably, Annex III of the AI Act subjects AI systems used in employment, education, and access to essential services to internal control-based conformity assessments that do not involve external certification bodies. As a result, high-risk workplace AI systems may be introduced and operated without meaningful external scrutiny.

While the regulation does introduce certain obligations for *deployers* – such as employers – these remain limited in scope. Employers must ensure that systems are used in accordance with the provider's specifications, and they are tasked with implementing human oversight and suspending use where risks to health, safety, or fundamental rights are identified. However, they are no longer generally required to conduct fundamental rights impact assessments, except in narrowly defined ipotesys. Furthermore, employers bear responsibility for communicating with trade unions and workers prior to the introduction of AI tools, providing information about the system's functions and objectives. This procedural transparency, while welcome, is insufficient on its own to guarantee substantive accountability – particularly when the underlying datasets and algorithms remain inaccessible and opaque.

Despite the AI Act's stated commitment to fundamental rights[16], the regulation does little to integrate the safeguards already present in data protection law. It does not, for instance, ensure that AI systems will be deployed in ways consistent with the GDPR's principles of necessity, proportionality, or fairness. Instead, by positioning AI systems within a distinct regulatory orbit, the AI Act may unintentionally marginalize the GDPR's protective logic – particularly its emphasis on limiting both the quantity and the scope of data collected.

The interdependence between the AI Act and the GDPR is undeniable,

[14] Peruzzi, *Intelligenza artificiale e diritto. Uno studio su poteri datoriali e tecniche di tutela*, Giappichelli, 2023.

[15] With specific reference to "high-risk" systems, the Regulation itself provides that the classification it establishes may be waived under certain conditions and based on a self-assessment conducted by the provider (Art. 6, par. 3 and 4).

[16] Tebano, *Intelligenza artificiale e datore di lavoro: scenari e regole*, in *DLM*, 2024, 3, p. 1 ff.

especially in contexts like employment, where personal data are continuously generated, processed, and evaluated. However, the AI Act fails to provide a coherent framework for reconciling its own risk-based regulatory model with the rights-based logic of data protection. The two systems risk operating in parallel rather than in synergy.

Against this backdrop, the next section turns to a detailed examination of the core tensions between AI deployment in the workplace and the application of data protection principles – focusing in particular on "purpose limitation" (§ 5) and "data minimization" (§ 6).

5. *The erosion of purpose limitation in the age of adaptive AI: from determinism to opacity*

The principle of "purpose limitation" is one of the foundational tenets of the GDPR. It mandates that personal data must be collected for specific, explicit, and legitimate purposes and must not be further processed in ways incompatible with those initial aims, unless a new legal basis is identified. Within the employment context, such a basis is typically found in the employer's "legitimate interest" (Article 6, par. 1(f), GDPR), but never in the consent of the employee – given the inherently unbalanced nature of the employment relationship.

The application of this principle is relatively straightforward in relation to deterministic AI systems – those whose outputs are predictable because they operate based on fixed, pre-programmed rules. In such cases, the employer, acting as "data controller", is required to clearly predefine the purposes of the data processing and ensure these purposes are transparent to the worker.

For example, consider a digital forensics tool (or "e-discovery" system) implemented to protect corporate assets pursuant to Article 4(1) of the Workers' Statute. If such a tool is deployed to automatically scan emails for keywords suggestive of illicit activity with a view to initiating disciplinary proceedings, it cannot subsequently be repurposed for a fundamentally different goal – such as quantitatively tracking employee email traffic to assess performance. Such a shift would constitute a violation of the principle of purpose limitation, unless grounded in a compatible legal basis and properly disclosed to the worker in advance.

This relatively clear framework begins to break down, however, when one considers "non-deterministic" or "adaptive" AI systems – those capable of learning from historical data and modifying their behavior over time without direct human intervention. These systems refine their outputs based on the patterns they detect, evolving continuously in both how they classify behavior and how they prioritize risk. As such, they may eventually requalify a worker's conduct based on new patterns, reclassify legitimate anomalies as suspicious, or even shift their internal thresholds for intervention in response to emerging correlations in data. What was once considered normal may later be flagged as deviant – not because of any actual change in conduct, but due to the model's evolving internal logic.

To illustrate, consider again an AI system deployed under the justification of protecting corporate assets. If this system is designed for cybersecurity purposes – such as an AI-driven threat detection platform – it may initially be calibrated to detect specific risk indicators, like keywords in emails. Over time, however, it may begin to flag actions such as transferring files from a different device, accessing records outside business hours, or logging in from a new location. While each of these behaviors may be entirely legitimate (e.g., due to remote work or workstation changes), the system's adaptive functioning may nevertheless classify them as suspicious.

The result is a form of surveillance that no longer targets specific, predefined conduct but instead operates through probabilistic profiling and open-ended anomaly detection. Workers, in turn, may be compelled to justify legitimate actions simply because they deviate from the system's expectations, thus experiencing a form of control that is both diffuse and opaque. This transforms the monitoring process from one of targeted oversight to one of continuous behavioral evaluation based on shifting and unpredictable criteria.

Moreover, the reliance on such systems places additional burdens on corporate IT personnel, who are now expected to continuously audit and recalibrate algorithmic outputs as part of their human oversight responsibilities. More fundamentally, however, adaptive AI challenges the very possibility of complying with the principle of purpose limitation: if the system's logic evolves, and if its operational focus shifts over time, how can the purposes of data processing be clearly defined and communicated in advance?

In this context, the legal obligation to provide clear and intelligible information to workers about how their data will be used becomes increasingly

difficult to fulfil. The dynamic nature of adaptive AI undermines the principle of foreseeability in data processing, revealing a structural incompatibility between the regulatory expectations of purpose specificity and the technical architecture of machine learning-based monitoring systems. This incompatibility demands urgent regulatory attention, particularly in the field of labor law, where the stakes for fundamental rights are especially high.

6. *Challenging the principle of data minimization in ai-driven workplaces*

Article 5(1)(c) of the GDPR enshrines the principle of "data minimization", requiring that personal data be "adequate, relevant, and limited to what is necessary in relation to the purposes for which they are processed." This obligation is reinforced by the technical principles of "privacy by design" and "by default", which demand that any technology employed by a data controller be configured to restrict data collection to the strict minimum required for achieving a predetermined, legitimate objective.

Beyond technical configuration, this principle also extends to the organizational dimension of data governance. Data controllers are encouraged to implement policies that favor targeted or randomized monitoring strategies over indiscriminate or continuous surveillance practices. In the employment context, this would entail preventive interventions oriented toward discouraging misconduct, rather than sustained, high-intensity tracking of individual behavior.

However, in contrast to the principle of purpose limitation – whose applicability depends to some extent on whether the AI system in question is deterministic or non-deterministic – the principle of data minimization is consistently undermined by AI systems across the board. Indeed, AI technologies, even those of moderate complexity, function optimally only when fed with large and diverse datasets. Their efficacy, and in some cases their very operation, presupposes a volume and granularity of data that is structurally at odds with the minimization imperative.

Illustrative examples abound. In the gig economy, platforms systematically collect and process workers' geolocation data, attendance records, acceptance or refusal of shifts, and delivery times in order to generate behavioral profiles and performance scores. In more traditional employment sectors, the widespread adoption of "fall detection" systems – typically based

on accelerometers embedded in wearable devices – entails the continuous monitoring of bodily movement patterns to trigger automatic alerts in the event of presumed physical distress. Similarly, the integration of "smart personal protective equipment" (PPE) for health monitoring – where the occupational physician acts as data controller – exemplifies a legal use of remote tracking that nonetheless challenges the boundaries of proportional data collection.

From a regulatory standpoint, the AI Act stipulates that any AI system, deterministic or otherwise, used for behavioral profiling in the workplace will automatically be classified as "high-risk" under Article 6(3)(d). While this classification reaffirms the requirement for compliance with the GDPR, it stops short of imposing a fully integrated normative framework on producers and deployers – one that would substantively guarantee compliance with the principles of purpose limitation and minimization. As such, the regulation remains largely procedural, leaving core normative tensions unresolved.

Responsibility for resolving these tensions, therefore, rests with the deployer, who acts as the data controller and is bound by the GDPR's accountability framework. This implies a proactive duty to identify, implement, and document technical and organizational measures aimed at reducing the risks associated with AI-driven data processing. These safeguards must be explicitly included in the "data protection impact assessment" (DPIA), and where compliance with the minimization principle cannot be reasonably assured – even through mitigation –profiling activities should not proceed. The mere operational value or perceived necessity of an AI system does not absolve employers from their legal obligation to protect the fundamental rights of workers.

Beyond the legal duty, the principle of data minimization should inform broader organizational decisions about the appropriateness of introducing AI tools in place of existing human supervision or simpler, non-adaptive technologies. Minimization must not be treated as a mere technical constraint but as a substantive ethical and legal consideration embedded within the corporate decision-making process itself.

Finally, it bears repeating that the enforceability of this principle would be significantly strengthened by embedding collective oversight mechanisms in workplace governance structures. Had the European regulatory framework mandated structured forms of worker consultation or codetermined

decision-making regarding the deployment of AI-based monitoring systems, the proportionality standard embodied in the data minimization principle would have been afforded a more effective and enforceable status.

7. *Automated decision-making in the workplace: limits of individual rights and the need for collective oversight*

The use of Artificial Intelligence (AI) in the workplace introduces significant challenges in terms of data protection and the regulation of decision-making processes. In particular, automated decision-making (ADM) systems – whether deterministic, non-deterministic, or generative – are increasingly used to manage tasks ranging from hiring to performance evaluation, scheduling, and even disciplinary measures. These systems process large amounts of data and generate outputs that can substantially affect the rights and freedoms of workers. While the General Data Protection Regulation (GDPR) and national labor law have introduced specific protections, a deeper analysis reveals the structural fragility of current safeguards, particularly when ADM systems are hybrid and when collective rights are neglected in favor of individual ones.

Under Article 22(1) of the GDPR, data subjects have "the right not to be subject to a decision based solely on automated processing, including profiling, which produces legal effects concerning him or her or similarly significantly affects him or her". Paragraph 3 of the same article requires that individuals subject to such decisions must be able to obtain human intervention, express their point of view, and contest the decision. Complementary obligations under Articles 13(2)(f), 14(2)(g), and 15 ensure ex-ante and ex-post transparency[17].

However, these safeguards are limited to decisions made solely by automated means, thereby excluding a broad spectrum of hybrid systems, where human oversight is nominal or merely formal. As scholarly literature highlights, the presence of a human "in the loop" does not necessarily mitigate the opacity or potential bias of algorithmic systems, especially when human operators lack the technical expertise to evaluate algorithmic outputs

[17] Wachter, Mittelstadt, Floridi, *Why a Right to Explanation of Automated decision-making Does not Exist in the General Data Protection Regulation*, in *IDPL*, 2017, 7, 2, p. 76 ff.

critically. Automation bias, where human decision-makers defer to algorithmic recommendations, often renders the human check ineffective[18].

Moreover, the exercise of individual rights under Article 22, such as access and contestation, proves largely inadequate in practice. Workers rarely possess the necessary information, time, or resources to interpret complex algorithmic logics or source code. Even when access is granted, trade secrets and intellectual property protections – recognized under Directive (EU) 2016/943 and reinforced by Recital 63 and Article 15(4) GDPR – often limit the disclosure of meaningful insights into algorithmic functioning[19].

Given these limitations, it is necessary to reconceptualize oversight not as an individual endeavor but as a collective right. Article 1-bis of Legislative Decree 152/1997, as amended by Legislative Decree 104/2022 and then by Legislative Decree 48/2023, partially addresses this issue by requiring employers to inform both individual workers and trade unions about the use of fully automated decision-making systems. However, by limiting the obligation to "fully" automated systems, the law enables circumvention where minimal human involvement is maintained.

A more effective solution would be to strengthen the role of trade unions by allowing them, with the aid of technical experts, to conduct independent audits of ADM systems. This should include access to technical documentation, training datasets, and, where appropriate, to portions of the source code – not to replicate or exploit the software, but to verify compliance with labor and data protection rights. Such oversight could be carried out under conditions that protect intellectual property and trade secrets, following the model of controlled access found in Article 22(3) GDPR and the GDPR-Recital 63 limitations.

In conclusion, hybrid ADM systems challenge the foundational assumptions of both data protection law and labor law. Individual rights are insufficient to counterbalance the algorithmic opacity and the systemic nature of decisions affecting workers. Therefore, the regulatory architecture must

[18] Kaminkski, Urban, *The Right to Contest AI*, in *CLR*, 2021, 121, p. 1957 ff.

[19] See Court of Justice of the European Union, 27 February 2025, CK, Case C-203/22, regarding the right of access under Article 15(1)(h) GDPR, clarified that the data subject is entitled to receive meaningful and comprehensible information about the actual logic applied in automated processing. This applies even when the information involves elements protected as trade secrets, the disclosure of which must be assessed by the competent supervisory authority or court through a balancing of the rights and interests at stake.

evolve to include stronger collective guarantees, more robust technical transparency mechanisms, and a structural rethinking of the human–machine relationship in employment contexts. Only by moving beyond individualistic paradigms can the law meaningfully respond to the challenges posed by artificial intelligence in the workplace.

Abstract

Artificial Intelligence is reshaping the workplace, but the legal framework designed to protect workers' privacy is struggling to keep pace. This article challenges the notion that AI requires exceptional legal treatment, arguing instead that it magnifies long-standing tensions within data protection and labor law. It highlights two core structural flaws: the illusion of individual control over personal data and the limits of accountability in algorithmic environments. As the AI Act introduces new rules, it risks sidelining key GDPR principles – such as purpose limitation and data minimization – by failing to confront the complexity of adaptive and non-deterministic systems. The paper focuses on the critical issue of automated decision-making, where existing safeguards, like Article 22 GDPR, often fall short. It calls for a shift from individual empowerment to collective oversight, empowering trade unions to scrutinize algorithmic systems – including their source code – while navigating the sensitive balance with trade secret protections.

Keywords

Artificial Intelligence, GDPR, Automated decision-making, Trade Union, Collective oversight.