

WORKFORCE DIGITAL TWINS IN EMPLOYMENT CONTEXTS

A Legal Analysis of Privacy, Consent, and Liability

Phan Ngoc Gia Bao

Faculty of Jurisprudence, Hanoi Law University, Hanoi, Vietnam

Email: phannngocgiabao123@gmail.com

Abstract

Purpose: *Workforce digital twins are ongoing, AI-based simulations of identifiable employees currently deployed on the platforms of HR giants like Microsoft, SAP, and Workday. Most of the companies using workforce digital twins and the vendors providing them are systematically in breach of Article 22 of the GDPR as well as their corresponding counterparts in the US. It is this proposition that will be developed in this paper, through doctrinal analysis of how current laws actually apply to such systems.*

Design/Methodology/Approach: *Doctrinal legal analysis is the main approach, since workforce digital twins are subject to a variety of laws which were not conceived as having to coexist with each other, and therefore need careful study in order to understand where there are inconsistencies, redundancies, and gaps in these regulatory systems. In this case, the focus will be on Articles 9 and 22 of the GDPR, particularly the recent CJEU interpretations in SCHUFA Holding (C-634/21, 2023) and D&B Austria (C-203/22, 2025), and the relevant sections of the Illinois Biometric Information Privacy Act, California Consumer Privacy Act, and EU AI Act's classification for high-risk systems*

Findings: *SCHUFA Holding (C-634/21, 2023) demonstrates that Article 22 comes into play whenever the output of the twin takes a decisive part in employment processes, despite any apparent human assessment step involved. The relationship between Article 22 and the EU AI Act's Article 14 human oversight obligations is not straightforwardly complementary. Article 22 gives the data subject a right to contest; Article 14 imposes on the operator a duty to design for meaningful oversight. The two provisions target different parties and require different things, and both apply simultaneously to the same deployment. US instruments, particularly the Illinois BIPA, impose biometric liability that is in several respects more acute than the GDPR equivalent. A four-part legal risk framework is developed, together with the implications of the May 2026 Digital Omnibus deferral of Annex III compliance to December 2027.*

Research Limitations/Implications: *Because no legislature drafted these instruments with workforce digital twins in mind, the analysis depends throughout on purposive interpretation -- and the practical consequence of that dependence is that two employers facing identical deployments could, in good faith and on reasonable legal advice, reach opposing conclusions about their compliance obligations under the same provision. Deployment rates and supervisory enforcement patterns in this sector are*

almost entirely undocumented; whether the interpretive positions taken here match how data protection authorities are actually handling complaints is, at this point, unknowable. Section 8's framework is a defensible starting point for compliance planning. It is not a prediction of outcomes.

Originality/Value: *Previous academic discourse has seen workforce digital twins in terms of a common challenge in governance rather than a new category that needs a specific approach in law. This approach fails to recognize a unique feature: workforce digital twins create a combination of obligations for continuous profiling, exposure of special category data, and product liability that is not created by any other use case scenario.*

Keywords: *algorithmic governance; automated decision-making; biometric data; employment law; EU AI Act; GDPR Article 22; human digital twin; worker data rights*

1. INTRODUCTION

The digital twin began as an engineering instrument, a real-time simulation counterpart to physical hardware first developed for NASA's Apollo programme. That origin is largely forgotten now. By 2032, the global digital twin market is projected to reach USD 259 billion, a figure powered by decades of convergence between artificial intelligence, machine learning, networked devices, and cloud infrastructure. Growth has not proceeded without anxiety. Kustelega et al. (2024), drawing on a 2021 Accenture survey, report that 73 per cent of respondents were already worried about digital twin privacy risks and 79 per cent of surveyed organisations identified GDPR compliance as a leading deployment concern. What is surprising, given that level of awareness, is how little legal scrutiny has been directed at what may be the most consequential application of the technology: the use of digital twins to model not infrastructure or machinery, but individual human employees.

This is not a speculative scenario. Major HR technology vendors already embed workforce digital twins as standard product offerings, not always using that term. Microsoft Viva Insights builds continuous behavioural profiles from employees' collaboration data; SAP SuccessFactors Workforce Analytics generates predictive scores for succession readiness and attrition risk; Workday People Analytics generates ranked outputs for promotion and performance decisions, and is used in redundancy and succession planning contexts, from live operational records. Taken together, these products are workforce digital twins in substance: persistent, probabilistic models of identified workers whose outputs shape decisions that matter to those workers' careers. The distinction from infrastructure or manufacturing deployments is not merely semantic. The subject of the twin is a person whose employment, income, and professional future may turn on what the model concludes about them, and who, in the vast majority of deployments, has no idea the model exists.

Governance scholarship on algorithmic employment has reached a broad consensus: the existing instruments fall short. Abraha's ILO Working Paper (2025) locates the root problem in structural

power asymmetries between employers and employees. Aloisi, Joppe, and Md Wasim Ahmed and Md Joshim Uddin, "Digital Twins and Legal Liability: Navigating Accountability in Simulated Realities" (2025) 7(8) *International Journal of Law and Politics Studies* 19, 19; Mihir Parikh, Aaron Kamath, Varsha Jajesh and Palak Kapoor, "The Tour d'Horizon of Data Law Implications of Digital Twins" (2025) 27 *OMG Journal of Innovation* 4 (Nishith Desai Associates).

Parikh et al. (n 1) 4, citing Fortune Business Insights Digital Twin Market Size Report (2024); Alexander Nechesov, Ivan Dorokhov and Jari Ruponen, "Virtual Cities: From Digital Twins to Autonomous AI Societies" (2025) 13 *IEEE Access* 13866, 13866-13867.

Marija Kustelega, Renata Mekovec and Ahmed Shareef, "Privacy and Security Challenges of the Digital Twin: Systematic Literature Review" (2024) 30(13) *Journal of Universal Computer Science* 1782, 1783, citing Accenture, "The Power of Massive, Intelligent, Digital Twins" (2021).

Microsoft, "Microsoft Viva Insights: Privacy, Security and Compliance" (Microsoft Technical Documentation, 2025); SAP, "SAP SuccessFactors Workforce Analytics: Product Overview" (SAP Documentation, 2025); Workday, "Workday People Analytics" (Workday Product Documentation, 2025).

Halefom Abraha, "Navigating Workers' Data Rights in the Digital Age: A Historical, Current, and Future Perspective on Workers' Data Protection" ILO Working Paper 149 (International Labour Office, Geneva, September 2025) 3-5.

Abraha (2025) press the analysis into individual instruments, showing that each of the GDPR, AI Act, and Platform Work Directive fails individually to provide context-specific protection for workers. Jørgensen and Ma (2025) map the specific transparency and accountability gaps those failures leave across six EU digital laws. That consensus is accurate as far as it goes. What it does not ask is whether existing instruments, purposively interpreted and applied together, might already demand more than current practice concedes. That is the question this paper takes up.

The analysis proceeds in stages, but the stages are not arbitrary. Sections 2 and 3 establish the analytical ground, covering what the existing literature has addressed, what it has not, and why the right starting point here is a doctrinal question rather than an empirical or normative one. In section 4, the notion of the workforce digital twin is first established as an object in law in order to consider under which legal framework it falls, as it turns out that its nature will have consequences for the issue at hand. In sections 5-7, the various frameworks applicable are then applied one by one: data protection, liability, and EU AI Act.

2. LITERATURE REVIEW

Four distinct bodies of scholarship engage with digital twin governance, and none was developed with

the workforce deployment context in mind. Healthcare and medical device governance, infrastructure and smart city applications, general data protection and privacy, and AI governance each raise questions relevant to the workforce case. The problem is that each body of literature approaches the twin as a different kind of object; the workforce digital twin sits where all four meet, and none of them addresses it from that position.

The healthcare digital twin literature has focused primarily on ethical and technical questions of deployment safety. Brand and Devlin (2025) catalogue data privacy, consent frameworks, model validation, and regulatory compliance as the dominant legal constraints in clinical settings. Gore et al. (2025) address the more specific problems of informed consent and medical device classification in biomedical and pharmaceutical deployments. Smart city and infrastructure applications raise a different set of questions. Peldon et al. (2024) examine data localisation and public interest governance as the defining challenges in urban twin deployments, while Nechesov, Dorokhov, and Ruponen (2025) address the governance of AI-driven twin systems operating at autonomous urban scale. What the healthcare and infrastructure streams share is an architecture of public or quasi-public accountability: clinical consent frameworks, public interest governance, and device classification all give affected persons some standing in the regulatory process. The employer-employee relationship provides

Antonio Aloisi, Anne Joppe and Halefom Abraha, "From 'General' to 'Context-Specific' Data Protection for Workers" (first published online 12 September 2025) *European Labour Law Journal*, DOI:10.1177/20319525251375031.

Bo Nørregaard Jørgensen and Zheng Grace Ma, "Digital Twins Under EU Law: A Unified Compliance Framework Across Smart Cities, Industry, Transportation, and Energy Systems" (2025) 14(24) *Electronics* 4881, DOI:10.3390/electronics14244881.

Dirk Brand and Sophia Devlin, "Ethical, Legal, and Technical Considerations for the Development of Digital Twins in Healthcare" (IEEE Conference Publication, August 2025), DOI:10.1109/ICTCS63349.2025.11212607.

Solomon I Gore and others, "Ethical and Legal Implications of Digital Twin Deployment in Biomedical and Pharmaceutical Domains" in *Accelerating Product Development Cycles with Digital Twins and IoT Integration* (IGI Global Scientific Publishing 2025) 491.

Daniele Peldon and others, "Navigating Urban Complexity: The Transformative Role of Digital Twins in Smart City Development" (2024) 111 *Sustainable Cities and Society* 105583.

none of that structure, and this literature cannot simply be borrowed for the workforce deployment context.

Compliance-by-design has become the expected minimum standard for digital twin governance, and the data protection literature explains why this was predictable. Kustelega, Mekovec, and Shareef

(2024) survey the field and identify algorithmic opacity and data quality deficiencies as the most consequential failure categories, which means the compliance problem is architectural before it is legal. Lennon, Julien, and Quin (2025) pursue that insight into practice, developing a methodology that treats GDPR obligations as design requirements embedded from the outset; retrospective compliance attempts, their analysis implies, are structurally unlikely to succeed. Jørgensen and Ma (2025) apply a similar analysis across the full EU digital regulatory stack and reach a conclusion that the individual instruments leave unresolved: transparency and accountability gaps in AI-driven twin components remain the hardest challenge, and the regulatory framework provides no mechanism for closing them from the outside. The present paper addresses what this body of work does not: how compliance-by-design is supposed to function when the organisation responsible for compliance is simultaneously the party whose operational interests the twin serves.

Aloisi, Joppe, and Abraha (2025) reach a conclusion that is both correct and insufficient: none of the GDPR, AI Act, or Platform Work Directive provides, individually, the context-specific framework that algorithmic management in employment settings requires. Correct, because each instrument was designed with a different primary problem in mind. Insufficient, because the implication that a better instrument is needed misses the possibility that these instruments, applied together and interpreted purposively, already demand more than current practice concedes, and that is the argument Sections 5 to 8 of this paper make. Abraha's ILO Working Paper (2025) charts structural deficiencies in worker data protection and makes the normative case for worker-centric governance frameworks; it is the right long-term prescription, but one that awaits the political conditions for implementation. Teller (2021) provides the conceptual footing for Section 4: the digital twin occupies an uncomfortable position between persons and things, carrying a patrimonial dimension that existing civil law categories have no ready means to accommodate. Ahmed and Uddin (2025) map the liability exposure this creates, identifying design defects, simulation defects, and data accuracy failures as distinct and cumulative heads. This paper draws on all of these, but its contribution is in the synthesis. The workforce digital twin cannot be properly analysed from within any single disciplinary frame.

Rosin and Parviainen (2025) examine the Platform Work Directive's transparency provisions in comparative context and find them wanting, though more charitably than the present analysis would. Their conclusion that the Directive moves the needle on transparency rights without

Kustelega, Mekovec and Shareef (n 3) 1789-1799.

Yann Lennon, Nathalie Julien and Annabel Quin, "Ensuring Personal Data Compliance by Integrating Legal Constraints into Digital Twin Design Methodology" in Proceedings of the 35th ESREL & 33rd SRA-Europe Conference (Research Publishing, Singapore 2025) 1618, 1619.

Jørgensen and Ma (n 7), where transparency and accountability gaps in AI-driven twin components are identified as the most intractable compliance challenge across the EU digital regulatory

framework.

Aloisi, Joppe and Abraha (n 6), on the analysis that none of the three instruments provides the context-specific framework the employment relationship demands when examined individually.

Abraha (n 5), mapping structural deficiencies in worker data protection across jurisdictions and proposing worker-centric data governance frameworks.

Marina Teller, "Legal Aspects Related to Digital Twin" (2021) 379(2207) *Philosophical Transactions of the Royal Society A* 20210023, 1-2.

Ahmed and Uddin (n 1), identifying design defects, simulation defects, and data accuracy defects as the three distinct heads of digital twin product liability.

resolving the structural power asymmetries of the employment relationship is accurate as far as it goes; what it understates is that the Directive was not designed to resolve those asymmetries and will not do so regardless of how it is implemented. Minotakis (2025), engaging the same two instruments from a more critical standpoint, draws on a 2025 OECD survey of over 6,000 firms whose results document both widespread algorithmic management adoption and tenacious accountability and transparency deficits.

What is missing from the existing literature is not a better instrument or a normative critique, but a combined doctrinal analysis that treats the workforce digital twin as a legal category in its own right. One area alone does not provide insight into the whole picture for compliance. Data protection doctrine helps to identify the issue regarding the legal basis and lack of consent; liability theory helps determine the result if the model itself falls apart; AI governance determines the standards for design. None of these frameworks was built with the others in mind, and each reaches a partial answer when applied in isolation. Applied together, they converge on a compliance architecture that is more demanding than any of them, read alone, suggests. Whether the workforce digital twin is sufficiently distinctive to warrant this combined analysis is an argument this paper makes rather than assumes. It turns on the legal character of the object itself, which is the question Section 4 takes up first.

3. METHODOLOGY

Doctrinal legal analysis is the method adopted here because the alternative approaches face a prior obstacle. Empirical analysis of compliance practice is largely blocked by the absence of publicly accessible DPIA records, vendor contract terms, and supervisory enforcement data in this sector. The Garante's November 2024 enforcement action against Foodinho S.r.l. for algorithmic management violations, resulting in a €5 million fine for, inter alia, failure to provide a compliant Article 22 human review mechanism, is among the few public reference points available. It is notable as an exception rather than a template for a well-documented enforcement landscape. Political economy analysis would be valuable but presupposes a settled account of what the existing law demands; that account is

what this paper attempts to provide. Doctrinal analysis is the necessary prior step.

The instruments selected converge on this deployment context in ways that a conventional regulatory survey would not capture. For EU-based deployments, the GDPR is the governing data protection framework, but its practical adequacy turns on the Article 22 scope and lawful basis questions that SCHUFA (C-634/21) and D&B Austria (C-203/22) have now given judicial content. US instruments are examined because the BIPA generates a qualitatively different liability profile from GDPR enforcement: workers constitute 78 per cent of settled BIPA cases in Edgeworth Economics' analysis of over 100 resolved actions, and the private cause of action, which requires no proof of actual harm after *Rosenbach v Six Flags* (2019),

Annika Rosin and Henni Parviainen, "Improving the Transparency of Algorithmic Management in Platform Work" (2025) *European Labour Law Journal* DOI:10.1177/20319525251375023.

Alexandros Minotakis, "Regulating AI in the Workplace: A Critique of the EU AI Act and the Platform Work Directive through a Worker-Centred Lens" (2025) *New Technology, Work and Employment* DOI:10.1177/29768624251396248.

Garante per la protezione dei dati personali, Provvedimento contro Foodinho S.r.l. (November 2024), imposing a €5 million fine for, inter alia, breach of GDPR Article 22 through an algorithmic management system governing shift allocation and de facto dismissals without a compliant human review mechanism; DataGuidance, "Italy: Garante Fines Foodinho EU5M for Processing Biometric Data" (November 2024) <<https://www.dataguidance.com/news/italy-garante-fines-foodinho-eu5m-processing-biometric>>. A prior enforcement action against the same entity attracted a €2.6 million sanction in 2021.

produces a scale of litigation exposure with no GDPR equivalent. What makes the EU AI Act necessary to examine here is not a further compliance checklist but an unresolved interaction: Article 14's structural oversight obligations apply simultaneously with GDPR Article 22 to the same deployment, and no regulatory guidance has yet addressed how those two frameworks operate together in HR analytics contexts. Civil liability is tackled through the updated Product Liability Directive (2024/2853). Should an employee find his/her career prospects negatively impacted due to being analyzed using a workforce twin based on data from training datasets that are discriminatory against him/her, or those records which do not represent who he/she is anymore, then there is a legal avenue open for such a claim.

Searches of Westlaw, HeinOnline, Google Scholar, and IEEE Xplore were conducted using terms including digital twin, workforce, algorithmic management, automated decision-making, and human digital twin combined with the names of the specific instruments examined. Publications from 2024 and 2025 were prioritised, though foundational earlier works are cited where the analysis requires. None of the instruments examined was drafted with workforce digital twin architectures in mind. The

analysis depends throughout on purposive interpretation, which carries a real risk: a supervisory authority, a court, or a regulator examining the same provisions may reach different conclusions, and some of the interpretive positions taken here have not yet been tested in enforcement proceedings.

4. THE WORKFORCE DIGITAL TWIN AS A LEGAL OBJECT

What distinguishes the workforce digital twin from other applications of the technology defined in ISO/IEC 30173:2023 (a synchronised, data-driven virtual representation operating at a specified frequency and fidelity relative to its physical counterpart) is that the physical counterpart is a person. Nechesov, Dorokhov, and Ruponen apply this standard's definition to persistent AI-driven simulations of complex human behavioural and performance systems; in the employment setting that scope matters because the simulation draws simultaneously on the full range of data an individual generates at work: structured HR records covering appraisals, compensation history, tenure, and absence; behavioural traces from enterprise collaboration platforms; operational output metrics including productivity logs and KPI scores; and, in more intensive deployments, physiological or biometric data gathered through keystroke dynamics, voice stress analysis, or wearable-integrated sensors. These are not instrument readings. They are a person's working life reconstructed in probabilistic form, and every new data ingestion event independently triggers fresh processing obligations under applicable data protection law.

Figure 1 provides an overview of the system architecture for a typical workforce digital twin implementation, starting from data collection, then passing through machine learning processing, and finally leading to employment decisions and major areas of legal liability.

DATA INPUTS
HR Records (appraisals, tenure, compensation, absence) Behavioural Data (email metadata, calendar, chat) Operational Metrics (output, KPIs, productivity logs) Biometric/Physiological Data (voice, keystroke, wearable) Third-Party Vendor Data (market benchmarks, peer comparisons)

Edgeworth Economics, "Analyzing Biometric Data Privacy Class Action Settlements" (Bloomberg Law database analysis, 2024), reporting that workers constitute 78 per cent of settled BIPA cases, with post-Cothron average per-class-member awards of \$1,049

Parikh et al. (n 1) 4, citing Digital Twin Consortium, "Definition of a Digital Twin" (official definition); ISO/IEC 30173:2023 Digital twin: Concepts and terminology (International Organization for Standardization 2023); Nechesov, Dorokhov and Ruponen (n 2) 13866-13868.

<i>Continuous ingestion</i>
ML PROCESSING LAYER - Supervised learning on historical workforce data; generates probabilistic scores per employee
<i>Twin outputs</i>
Flight Risk Score Redundancy Priority Rank Succession Readiness Performance Trajectory Promotion Candidate Score
<i>Nominal human review</i>
EMPLOYMENT DECISIONS: Promotion - Redundancy selection - Succession planning - Performance management (reviewer presented with ranked outputs; departure from recommendation requires justification)
LEGAL EXPOSURE: GDPR Arts 4, 9, 22 - Illinois BIPA - CCPA/CPRA - EU AI Act Annex III Cat. 4 - EU Product Liability Dir 2024/2853 Case law: SCHUFA C-634/21 (CJEU 2023); D&B C-203/22 (CJEU Feb 2025)

Figure 1: Architecture of a Workforce Digital Twin Deployment (Source: Author's representation based on vendor documentation and Lennon, Julien and Quin (2025))

Persistence changes the legal picture fundamentally. A periodic appraisal or a one-off redundancy scoring exercise is a snapshot; the twin is not. It is a continuously evolving model that updates in real time as new data flows in. Bolyn and Bechet identify this continuity and granularity as the feature that separates a digital twin from a conventional analytics tool, and its legal significance is what follows from it. Both Lennon, Julien, and Quin and Jørgensen and Ma note the same implication: each new data input independently triggers fresh processing obligations. The employee is never simply processed once and done. Profiling is ongoing, and every ingestion event carries its own requirements of consent, lawfulness, and proportionality.

Classifying a workforce digital twin as a legal object turns out to be a genuinely difficult question, and not merely in a technical doctrinal sense. Teller frames the conceptual difficulty as a tension between being and having: the twin simultaneously replicates a person and commodifies their data, generating a patrimonial dimension for which civil law has no ready-made category. The distinction between the digital representation and the human being becomes blurred in ways that force the law to evolve definitions of something that is neither quite human nor merely an object. The theory, advanced prior to the present widespread use of digital twins in the workplace, rings true to me in every respect. It is not simply an issue of classification; rather, it has real-world implications regarding who is liable for what.

In workforce deployments, the stakeholder topology of digital twin liability maps onto three distinct parties: the employee whose profile the twin embodies, the HR technology vendor who built and

operates the system, and the employing organisation that procures and deploys it. The critical structural feature of this setting, as Abraha emphasises, is that the employing organisation occupies two of these positions simultaneously: it is both the data controller and the party whose operational interests the twin's outputs are designed to serve. No other domain

Alex Bolyn and Eric Bechet, "Identifying a Problem and Solving It with a Digital Twin Idea" (Digital Twin Academy Module, Interreg V-A Euregio Meuse-Rhine, April 2023) 2, 8.

Lennon, Julien and Quin (n 12) 1621-1622; Jørgensen and Ma (n 7) 4881, on the point that each new data input independently triggers fresh processing obligations under applicable data protection law.

Teller (n 16) 1-2.

Abraha (n 5) 8-9; Teller (n 16) 2-3, identifying the coincidence of controller and beneficiary roles as the defining structural feature of human digital twin liability in employment contexts.

in which digital twins are currently deployed presents this structural conflict. In this regard, the discussion by BUSE concerning human digital twins highlights that these are particularly placed in an awkward juncture between privacy laws and personality rights and would need informed consent, appropriate contracts, and control measures as bare minimum.

5. DATA PROTECTION FRAMEWORK

5.1 GDPR Profiling, Special Category Data, and Lawful Basis

Most of the organizations using the workforce digital twin model seem to regard GDPR as a compliance checklist as opposed to a real regulation, and its provisions that apply to such an architecture go a long way to explain why this approach cannot last. The definition of personal data provided for in Article 4(1) as any data relating to a living individual who can be identified, directly or indirectly, includes not just the data that is inputted into the digital twin model but even the outputted data. Thus, a probabilistic redundancy score falls under the category of personal data, and calculating the redundancy score amounts to profiling as defined in Article 4(4) of the GDPR as automated processing in connection with the evaluation of certain aspects of the individual in their working life. Lennon, Julien, and Quin highlight another important aspect of the GDPR in this regard: its application goes beyond data which relates directly to named individuals.

The applicability of Article 9 to the stream of data inputted to a digital twin of the average working force is not a complex matter. Data concerning uniquely identifying the natural person via biometric means clearly constitutes processing of data falling under Article 9; for instance, keystroke and voice stress data are both biometric data as defined by the Article. Another area of interest in this context lies in the realm of health data: any conclusions drawn using the data of sensors worn on the body of employees or even absence data is covered by the term "health data" regardless of whether this is the system's characterization of the data. The heightened liability for special category data does not stop at

the data controller's boundary. It follows the processing wherever it goes, including to a contracted vendor; Parikh et al. and Abraha both document this.

Choosing the wrong lawful basis under Article 6 is not merely a technical error; for a deployment of this scope it exposes the controller to the risk that every processing activity built on that basis is unlawful *ab initio*. In practice, most employers appear to invoke either Article 6(1)(b), which permits processing necessary for performance of a contract, or Article 6(1)(f), the legitimate interests basis. Neither is adequate for the full scope of a workforce digital twin deployment, and the inadequacy is not a close question. Article 6(1)(b) requires that processing be necessary for the performance of a contract. Continuous probabilistic simulation of a worker's future career trajectory does not satisfy that test; no reasonable reading of an

Dirk Helbing and Javier Argota Sanchez-Vaquerizo, "Digital Twins: Potentials, Ethical Issues, and Limitations" [revised final version] (ETH Zurich / Complexity Science Hub Vienna 2022) 9; Abraha (n 5) 8-9; BUSE Attorneys at Law, "Digital Twins of Real People: Legal Risks for Businesses" (BUSE Law Blog, March 2026) <<https://buse.de/en/blog-en/labor-law/data-protection/digital-twins-of-real-people-legal-risks-for-businesses/>>.

Lennon, Julien and Quin (n 12) 1621-1622, distinguishing "explicit" personal data from data that can be combined to infer personal characteristics, both of which attract full GDPR obligations.

Teller (n 16) 2; Regulation (EU) 2016/679 (General Data Protection Regulation) [2016] OJ L119/1, art 9; BUSE Attorneys at Law (n 26), on Article 9(2)(a) GDPR consent as the only practically available lawful basis for biometric identification processing in employment contexts.

Parikh et al. (n 1) 14; Abraha (n 5) 18-20, on heightened liability for special category data following the processing wherever it goes, including to contracted processors.

employment contract makes it necessary. The legitimate interests basis requires a proportionality balancing exercise, and Aloisi, Joppe, and Abraha observe that the structural power imbalance inherent in employment fundamentally compromises both that balance and any reliance on consent in AI-driven workforce management contexts.

5.2 GDPR Article 22 and the Nominal Human Review Problem

Redundancy priority rankings, promotion scores, succession probabilities: these are the outputs workforce digital twins generate, and each is a decision that Article 22 covers: automated processing producing a legally significant or similarly significant impact on the worker. That the processing satisfies the threshold is not seriously contested. What has occupied most compliance attention is a different question: whether placing a nominal human reviewer between the automated output and the final employment decision takes the overall process outside Article 22's scope. That has been the standard mechanism by which deploying organisations have sought to avoid the provision's obligations, and the CJEU has addressed it directly.

In SCHUFA Holding (C-634/21, December 2023), the Court held that an automated probability score engages Article 22 wherever it plays a determining role in an adverse decision, even if the formal act of deciding rests with a third party, provided the score in practice leads to that outcome in almost all cases. The controller cannot bypass the provisions of Article 22 through a mere formality of having a human review process which makes no substantial difference at all. The court ruling of February 2025 on D&B Austria (C-203/22) made the discussion more precise in that it directly dealt with the conflict between the rights of the data subject and the score provider's interest in maintaining the model's confidentiality. The Court's resolution was that the controller must disclose relevant information to the competent supervisory authority or court for proportionality balancing; blanket withholding from the data subject is impermissible.

What the SCHUFA framework means for workforce digital twin deployments is worth stating plainly. Take the situation in which a line manager receives a twin-generated redundancy priority ranking with confidence scores and performance decile placements, and the organisational process treats those outputs as the baseline from which departure needs justification. The human review is there on paper. But the reviewer cannot interrogate why the model ranked one employee above another, has no access to the training data, and works in an institution where disagreeing with the system requires more justification than agreeing with it. Under SCHUFA, that reviewer is not the Article 22 decision-maker. The twin is

Contestability, understood as the employee's ability to interrogate the twin's assessment and have that interrogation change the outcome, is the structural requirement that transparency-focused compliance tends to understate, as Teller has argued. The point is not merely academic. Article 22's human involvement requirement is substantively empty when the

Kustelega, Mekovec and Shareef (n 3) 1783, 1792; GDPR (n 28) art 6; Lennon, Julien and Quin (n 12) 1621-1622; Aloisi, Joppe and Abraha (n 6), on the employment power imbalance "radically undermines" the consent and legitimate interest bases in AI-driven workforce management contexts.

Case C-634/21, SCHUFA Holding and Others (Scoring), Judgment of the Court of Justice of the European Union of 7 December 2023, ECLI:EU:C:2023:957, paras 47-48, 61, 73.

Case C-203/22, D&B Austria GmbH v CK, Judgment of the Court of Justice of the European Union of 27 February 2025, ECLI:EU:C:2025:117; Ilke Okan, "AI Gets Personal: CCPA vs GDPR on Automated Decision-Making" (Berkeley Technology Law Journal, April 2025) <<https://btlj.org/2025/04/ccpa-vs-gdpr-on-automated-decision-making/>>.

Teller (n 16) 4.

reviewer possesses neither the information nor the institutional standing to depart from the system's recommendation; a nominal challenge right that cannot in practice be exercised is not a challenge right at all. Helbing and Argota Sanchez-Vaquerizo situate this within a broader concern about

deterministic algorithmic systems eroding the conditions of individual agency that the legal order presupposes, an argument that sits uneasily with the AI Act's aspiration to legislate genuine human oversight into existence. According to Wagg et al., assurance, validation, and verification as part and parcel of a legitimate twin constitute a basic necessity and not a desirability in itself, and failure to meet any of these criteria renders such a workforce twin ineffective.

The Article 29 Working Party, whose guidance the EDPB has continued to apply, stated the principle clearly: a review devoid of genuine decisional authority does not satisfy Article 22's human involvement requirement. SCHUFA and D&B Austria have since turned that principle into binding CJEU authority. Article 22 compliance cannot, in consequence, be achieved by adding a nominal sign-off step to a process that was automated throughout.

5.3 The CCPA, CPRA, and the Illinois BIPA

For organisations deploying workforce digital twins in the United States, the regulatory challenge is not a single framework but an accumulating stack of state-level instruments with overlapping requirements, no federal floor, and in the case of the Illinois BIPA, no statutory cap on aggregate liability. The CCPA, as amended by the CPRA with effect from 1 January 2023, brought employees of covered California businesses within its scope. The CPPA's proposed ADMT rules, still pending finalisation at the time of writing, would require employers to notify employees before using automated technology as a substantial factor in consequential decisions, and would permit employees to opt out. The divergence in these proposals is examined by Okan (2025), in relation to the GDPR approach, whereby the right to automated decision-making provided by the GDPR mandates an opt-in process with express consent whereas the California model allows only an opt-out right, which does not apply to high-stakes employment applications such as recruiting.

The greatest legal risks, in the immediate future, to organisations that adopt workforce digital twins incorporating biometrics, are found under the Illinois Biometric Information Privacy Act 2008. The statute conditions any collection of biometric identifiers on prior written consent, with documented retention and destruction schedules as a further mandatory element. Commercialisation of biometric data is prohibited entirely. Its most consequential feature is a private cause of action with statutory damages accumulating rapidly in workforce contexts: USD 1,000 per negligent violation and USD 5,000 per intentional or reckless violation, assessed per person and per individual occurrence. In the case of *Cothron v White Castle System Inc* (2023), the Supreme Court of Illinois held that each scan/transmission is considered a discrete violation, rather than just an ingredient for a larger offense. Abraha's analysis of continuous AI-driven biometric monitoring concludes that cumulative BIPA exposure,

GDPR (n 28) art 22; Article 29 Working Party (n 37) 21–22, on a review devoid of genuine decisional authority being incapable of satisfying Article 22's human involvement requirement.

The greatest legal risks, in the immediate future, to organisations that adopt workforce digital twins incorporating biometrics, are found under the Illinois Biometric Information Privacy Act 2008. The statute conditions any collection of biometric identifiers on prior written consent, with documented retention and destruction schedules as a further mandatory element. Commercialisation of biometric data is prohibited entirely. Its most consequential feature is a private cause of action with statutory damages accumulating rapidly in workforce contexts: USD 1,000 per negligent violation and USD 5,000 per intentional or reckless violation, assessed per person and per individual occurrence. In the case of *Cothron v White Castle System Inc* (2023), the Supreme Court of Illinois held that each scan/transmission is considered a discrete violation, rather than just an ingredient for a larger offense. Abrahama's analysis of continuous AI-driven biometric monitoring concludes that cumulative BIPA exposure,

GDPR (n 28) art 22; Article 29 Working Party (n 37) 21–22, on a review devoid of genuine decisional authority being incapable of satisfying Article 22's human involvement requirement.

Helbing and Argota Sanchez-Vaquerizo (n 26) 11–12.

David J Wagg, Christopher Burr, Jason Shepherd, Zack Xuereb Conti, Mark Enzer and Steven Niederer, "The Philosophical Foundations of Digital Twinning" (2025) 6 Data-Centric Engineering e12, 14–15.

Article 29 Working Party, "Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679" (WP 251rev.01, 6 February 2018) 21–22; SCHUFA (n 31) para 73.

Illinois Biometric Information Privacy Act, 740 ILCS 14/1 et seq (2008), ss 20(1)–(2); Parikh et al. (n 1) 13.

Cothron v White Castle System Inc, 2023 IL 128004 (Supreme Court of Illinois, 2023).

calculated across all employees and the duration of deployment, is capable of reaching levels that threaten the financial viability of the deploying organisation.

Table 1 provides a comparative overview of the automated decision-making regimes applicable to workforce digital twin deployments across the jurisdictions examined in this section.

***Table 1: Comparative Overview of Automated Decision-Making Regimes
Applicable to Workforce Digital Twins***

Criterion	GDPR Art. 22 (EU)	CCPA/CPRA (California)	Illinois BIPA	EU AI Act Art. 14
Consent model	Opt-in; explicit consent	Opt-out (limits for high-impact uses)	Written opt-in per biometric	Pre-deployment notice and consultation
Human oversight	Substantive review required (SCHUFA)	Not mandated	Not addressed	Structural obligation under Art. 14

Private action	Via DPA; limited individual claims	CPPA enforcement	USD 1,000 to 5,000 per violation	Fines up to 7% annual turnover
Transparency	Meaningful information; D&B Austria limits trade-secret shield	Pre-use notice (proposed ADMT rules)	Disclosure of retention and destruction schedule	Technical documentation; DPIA; post-market monitoring

Note. ADMT = automated decision-making technology. CPPA = California Privacy Protection Agency. DPA = Data Protection Authority.

6. LIABILITY FRAMEWORKS

6.1 Product Liability and Negligence

The Revised Directive on Product Liability of the European Union (2024/2853) extends the coverage of such directives to software and artificial intelligence systems through the inclusion of "workforce analytics platforms". While there are several grounds of liability identified by Ahmed and Uddin with respect to digital twins, the design defect is the most persistent when it comes to the employment context. If the algorithms and models underlying the digital twin are based on a dataset that contains a significant under-representation of certain protected characteristics in the sample, then the result will be biased regardless of any review processes conducted after. The second ground of liability, however, is less obvious but not less important.

Abraha (n 5) 27-28, on per-violation aggregate BIPA liability generated by continuous AI-driven workplace biometric monitoring. Note: Illinois amended BIPA in August 2024 (SB 2979) to provide that repeated collections of biometric data from the same person using the same method constitute a single violation, limiting aggregate exposure relative to the position established in Cothron. The amendment applies retroactively (Seventh Circuit, April 2026). The liability remains substantial for new deployments without written consent.

As the statistical parameters of the model start drifting away from the actual values for the relevant workforces, predictions and recommendations become unreliable, but no one is going to question their validity in an organizational environment that assumes the results' accuracy. Thirdly, with regard to data accuracy, this concern is more readily addressed, but can have the immediate and adverse effect of producing a representation that fails to capture the actual individual. As shown by the systematic literature review by Kustelega et al. on digital twin applications, data quality was cited as one of the most common sources of failure in the industry. In the employment sector, in particular, failure to correct inaccurate data can result in job loss.

Negligence Liability poses more complex causation issues. According to Ahmed and Uddin, when courts deal with digital twin cases, there is a problem of structure as the reasoning of the digital twin is not open for examination by either the claimant or an expert witness. This is referred to as black box behaviour that characterizes machine learning where the path from input to output is beyond the reach of anyone who does not control the algorithm. The reason for making the data protection and AI law

principles relevant to torts is to provide the necessary evidence conditions.

6.2 Contractual Liability, Vendor Relations, and the Trade Secrets Problem

The usual pattern for the employment of a workforce digital twin entails a triangular contractual arrangement, whereby the employer purchases the technology from an HR tech company and commissions the company to analyze employees' data. Under GDPR, the employer holds the position of the data controller and the tech company that of the data processor. It means that the Article 28 contract should be obligatory before processing can start, and not vice versa. However, agreements that eventually come about fail to meet this benchmark. First, vendors are reluctant to agree on terms that will allow an external evaluation of their architecture; secondly, employers lacking technical expertise have little chance to negotiate; thirdly, contracts are written to the advantage of the vendors' business interests, not employee rights.

The weaknesses of the contract are described by Parikh et al. and Abraha in similar terms. Repeated weaknesses include any contractual provision regarding periodic bias audit and the transmission of results to the organization deploying the software, as well as any notice obligations in case of changes in model structure or training datasets; explainability outputs needed to give employees the possibility of challenging decisions under Article 22 of the GDPR; and indemnification provisions covering the risk of claims due to discriminatory outcomes from model failures or biased training data. Parikh et al. further suggest that the existing cyber insurance and professional indemnity policies should be analyzed for their adequacy vis-à-vis the risks involved in automated decision-making. The decision in D&B Austria (C-203/22, 2025) fundamentally undermines the opacity argument which is relied upon in many of today's contracts. A vendor cannot rely anymore on a simple invocation of trade secret protection as a shield against disclosure of how the model works. What the decision gives in its stead is proportional supervised disclosure.

7. THE EU AI ACT AND THE HIGH-RISK CLASSIFICATION OF WORKFORCE DIGITAL TWINS

Kusteleaga, Mekovec and Shareef (n 3) 1794-1795.

Ahmed and Uddin (n 1) 23; D&B Austria C-203/22 (n 32), establishing that the tension between vendor trade secret protection and the data subject's right to meaningful information must be resolved through supervised disclosure to a competent authority rather than blanket withholding.

The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) entered into force on 1 August 2024. Annex III, Category 4 designates as high-risk those AI systems intended for making decisions affecting the terms and conditions of work relationships, or for monitoring and evaluating the performance and behaviour of persons in those relationships. Any workforce digital twin generating outputs that feed into promotion, redundancy, succession, or performance management decisions falls

within this definition. The more difficult question, one that Aloisi, Joppe, and Abraha (2025) and Nechesov, Dorokhov, and Ruponen both stop short of fully addressing, is how the Annex III classification sits alongside the GDPR's Article 22 framework when both apply simultaneously to the same system. That co-application is the normal case, and the obligations are not identical.

What Articles 9 to 17 collectively demand, read together rather than in isolation, is a workforce digital twin whose architecture was designed for compliance before it was ever deployed. The most demanding single requirement is representativeness in training data: a model calibrated on historically stratified workforce records encodes that stratification in its outputs from day one, and no post-hoc audit corrects this. The most telling requirement is representativeness in training data, because it cannot be satisfied by documentation, a model built on stratified workforce records encodes that stratification from day one, and no governance layer added afterwards corrects it. The other obligations, continuous risk management, event logging for post-hoc review, structural oversight, pre-deployment conformity assessment, share the same characteristic. None of them can be retrofitted. They presuppose a system that was designed with compliance as an architectural goal. What the evidence from current commercial deployments suggests is that most systems were not designed that way.

The human oversight obligations under Article 14 are where the Act's demands become most architecturally disruptive. The requirement is that high-risk AI systems be designed such that assigned oversight persons can genuinely monitor the system's operation. Effective oversight also presupposes that those persons remain alert to the tendency to over-rely on algorithmic output, a tendency that intensifies precisely where the system's recommendations carry institutional authority, and that they are equipped to interpret what the output actually means in the specific decision context before them. The operative word is designed: the obligation is not satisfied by placing a human in the sign-off chain if the system was built in a way that makes meaningful review practically impossible. Helbing and Argota Sanchez-Vaquerizo argue that transparency must extend beyond technical documentation to a form accessible to affected persons; Jørgensen and Ma's analysis of the EU digital regulatory stack finds that explainability and accountability gaps in AI-driven twin components remain the most intractable compliance challenge. Both analyses arrive at the same problem from different directions. The systems currently deployed in workforce settings do not satisfy Article 14 as a matter of design. Adding audit logs and reviewer training does not change this.

A significant development occurred after the primary analysis in this paper was completed. On 7 May 2026, the European Parliament, the Council of the EU, and the European Commission reached a provisional political agreement on the Digital Omnibus on AI, deferring the high-risk AI compliance deadline for standalone Annex III systems from 2 August 2026 to 2

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial

Intelligence Act) [2024] OJ L2024/1689, Annex III Category 4, arts 9-17; Lennon, Julien and Quin (n 12) 1619; Aloisi, Joppe and Abraha (n 6); Nechesov, Dorokhov and Ruponen (n 2) 13870-13875.

Helbing and Argota Sanchez-Vaquerizo (n 26) 9, 15; EU AI Act (n 49) art 13; Jørgensen and Ma (n 7), on transparency and explainability gaps in AI-driven twin components as the most persistent structural challenge across the EU digital regulatory framework.

December 2027. The deferral reflects practical difficulties in designating national competent authorities and the absence of harmonised standards, not any relaxation of the underlying obligations. The substantive architecture of the AI Act is unchanged. The Omnibus also reinforces Article 26(7), which requires employers to inform workers and their representatives before deploying high-risk AI systems, a requirement that sits uneasily with the opacity that characterises most current deployments.

Wagg et al. identify the ethical, legal, and societal issues arising from digital twins used to represent living entities as a critical area for future research, noting that such twins create new capacities to influence decisions about human subjects in ways that generate normative challenges requiring systematic analysis. The EU AI Act provides the first legislative response to that challenge in the employment context. The response is structural. It does not merely require that a system be able to explain itself; it requires that systems be built so that genuine explanation is possible by design. That is a harder constraint, and the gap between what Article 14 requires and what current deployments actually provide is a design gap, not a paperwork gap.

8. A LEGAL RISK FRAMEWORK FOR WORKFORCE DIGITAL TWINS

Any organisation deploying a workforce digital twin faces legal exposure across four domains, and the domains interlock in ways that make a piecemeal approach structurally dangerous. Getting the data classification wrong forecloses the best available lawful bases; a weak lawful basis creates gaps in the Article 22 governance architecture; and an inadequate governance architecture leaves the vendor contract with nothing meaningful to demand of the processor. Table 2 maps these obligations. The subsections that follow explain why each layer depends on the one before it.

Table 2: Legal Risk Framework for Workforce Digital Twin Deployments

Dimension	Core Legal Obligations	Primary Instruments
1. Data Mapping and Classification	Map all input streams; classify personal, special category, and biometric data; conduct DPIA before deployment and on each new data input	GDPR Arts 4, 9, 35; ISO/IEC 30173:2023
2. Lawful Basis and Consent Architecture	Establish granular, purpose-specific lawful basis; obtain written biometric consent; implement meaningful withdrawal mechanism; adopt erasing policy	GDPR Arts 6, 7, 9; CCPA/CPRA; Illinois BIPA s 15

3. Automated Decision-Making Governance	Provide substantive (not nominal) human oversight; disclose twin existence and outputs to employee; enable meaningful	GDPR Art 22, Recital 71; SCHUFA C-634/21; D&B C-
---	---	--

8.1 Data Mapping and Classification

Getting the data classification wrong at the outset vitiates the lawful basis analysis for every processing step that follows, which is why this exercise cannot be deferred until a complaint arrives. Each input stream feeding the twin carries its own legal character: biometric data, special category health data, and ordinary HR records attract materially different obligations, and a deployment that treats them uniformly will be non-compliant in ways that may not surface until a supervisory authority intervenes. Lennon, Julien, and Quin's compliance-by-design methodology addresses this sequencing problem: characterise every data type in a Configuration phase, conduct a DPIA in a Design phase, and repeat both stages when new data is introduced. The DPIA question is less settled than it appears. Jørgensen and Ma's analysis of the EU digital regulatory stack shows that transparency and accountability gaps concentrate in deployments where design-stage governance decisions were deferred; their analysis suggests that organisations which do not resolve these questions before going live tend not to resolve them at all. Some national supervisory authorities apply the Article 35(3)(b) automated decision-making trigger narrowly, and the AI Act's Annex III Category 4 classification does not automatically settle the Article 35 threshold. That determination belongs with the relevant supervisory authority before deployment, not after a complaint is filed.

8.2 Lawful Basis and Consent Architecture

Each processing purpose within the twin architecture requires an identified lawful basis that is defensible in the employment context. The two most commonly invoked, contractual necessity under Article 6(1)(b) and legitimate interests under Article 6(1)(f), are not reliably available across the full scope of a workforce twin deployment, as the analysis in Section 5.1 demonstrates. The practical implication is that an organisation without an adequate lawful basis cannot lawfully process, and cannot be brought into compliance through remedial documentation alone. Where the organisation falls back on consent, the deeper difficulty is the withdrawal mechanism. An employee who withdraws consent to biometric processing mid-deployment does not simply halt future processing; the organisation must address what happens to the twin model that has already incorporated that employee's data and to the outputs it has already generated. This is not a question the GDPR resolves in terms specific to continuous processing architectures, and neither Abraha nor Aloisi, Joppe, and Abraha provide a worked answer. For biometric data in Illinois, the BIPA's requirement of a documented destruction

Parikh et al. (n 1) 22-23; GDPR (n 28) art 35; Lennon, Julien and Quin (n 12) 1621, 1623-1624; Jørgensen and Ma (n 7), on compliance-by-design as the minimum standard across the EU digital regulatory stack.

schedule at least forces the organisation to address the question in advance, which is rather the point of it.

8.3 Automated Decision-Making Governance

In D&B Austria (C-203/22, February 2025), the CJEU confronted a vendor's refusal to disclose how its automated scoring model reached its outputs: the data subject asked, the vendor cited trade secret protection, and the Court ruled that blanket withholding is impermissible. Supervised disclosure to a competent authority or court is the mechanism; unilateral refusal is not an option. That ruling built on what SCHUFA Holding (C-634/21, December 2023) had established: that an automated score engages Article 22 wherever it plays a determining role in a consequential decision, regardless of whether a human formally signs off at the end of the process. Between them, these two cases foreclose the nominal human review defence and the vendor opacity defence. The D&B Austria and SCHUFA holdings remove both scaffolding assumptions on which current deployments typically rest.

Institutions do not work the way legal frameworks assume they do. In a large corporate organisation, a redundancy decision made with reference to a twin-generated priority ranking is one item in a volume process managed by a reviewer who has no access to the model's uncertainty measures, no visibility into how demographic subgroups performed differently in the training data, and no organisational standing to depart from the output without written justification. Technical provision of disaggregated data does not solve this problem. The problem is that the institution treats the twin's recommendation as a baseline rather than as one input among several. Compliance with Article 22, properly understood, requires redesigning the decision process so that departure from the model is the normal case, not the exceptional one. On top of that, employees have rights under the SCHUFA framework that current practice ignores: the right to know the twin contributed to the decision, to know which factors were weighted, and to contest the assessment before any outcome takes effect. The EU AI Act's Article 14 sets the structural minimum for all of this. Almost no current deployment meets it.

8.4 Vendor Contracting and Accountability

The starting point for vendor contracting changed materially with D&B Austria (C-203/22, 2025). The vendor trade secret defence no longer operates as a blanket shield: information about model operation cannot be withheld entirely from competent authorities or courts, and the mechanism for handling the tension between confidentiality and disclosure is supervised proportionality balancing, not categorical refusal. That shift changes the negotiating dynamic. A vendor who can no longer credibly refuse all disclosure of model architecture has less leverage to resist contractual audit rights or explainability provisions. The Article 28 data processing agreement needs provisions that standard-

form contracts in this market do not include. Bias audits, disaggregated by demographic subgroup and conducted periodically, with results disclosed to the deploying organisation, are the most technically demanding. Advance notification before any material change to model architecture or training data matters because a compliant deployment can become non-compliant overnight if the vendor retrains without warning. Indemnity provisions should allocate liability for discrimination claims arising from model defects to the party best placed to prevent them. Employees exercising Article 22.

Lennon, Julien and Quin (n 12) 1622-1623; Abraha (n 5) 30-32, on employment power asymmetry as the structural obstacle to genuine consent in AI-driven workforce data processing.

challenge rights need explainability outputs they can actually use. Standard-form commercial agreements in this market fall well short of these requirements, and the gap is not incidental.

9. CONCLUSION

Workforce digital twins circulate today under commercial labels such as workforce analytics, people intelligence, and talent succession modelling. The legal governance of these systems has not kept pace. What this analysis shows is not a situation of clear rules inadequately enforced. The data protection instruments reach this deployment context but were calibrated for less continuous processing. The liability doctrine was not designed to penetrate machine learning opacity. The employment law instruments predate the concept of running a persistent probabilistic model of a named employee's working life. Three bodies of law, three structural gaps, and none of the gaps is remedied by the others.

That patchwork is more demanding than organisations typically acknowledge, and the routes around it have been closing. Nominal human review no longer satisfies Article 22 after SCHUFA and D&B Austria; vendor opacity no longer shields model logic from compelled disclosure to competent authorities. The BIPA problem is less a compliance question than an actuarial one: at USD 5,000 per intentional violation assessed per scan and per person, the arithmetic of non-compliance across a workforce with continuous biometric processing is potentially uninsurable, not merely serious. The AI Act adds a structural constraint rather than a financial one; it requires governance infrastructure that must exist before deployment, not as retrospective remediation. The CJEU has foreclosed nominal human review and vendor opacity. The BIPA has quantified the cost of operating without biometric consent. The EU AI Act has imposed a design standard that cannot be satisfied retroactively. Each instrument targets a different aspect of the same deployment. Together they leave very little room. Abraha (2025) and Aloisi, Joppe, and Abraha (2025) document this pattern across jurisdictions: the persistent problem is the exploitation of interpretive ambiguity and technical opacity, not the absence of applicable law.

Teller's conceptual framework of an entity halfway between people and things and Helbing and Argota Sanchez-Vaquerizo's analysis of the quantification tendency operate at different levels and do not resolve into a clean synthesis. Teller's concern is with classification: what legal category does the twin occupy, and where does liability sit? Helbing's concern is with epistemic loss: what happens to the dimensions of a person that a model cannot capture? In the workforce context, these two problems interact. Because the twin is neither fully a legal subject nor clearly a product, accountability for its outputs is diffuse. Because it can only score what is measurable, the decisions it informs are systematically skewed against whatever aspects of an employee's contribution resist quantification. Wagg et al.'s insistence on assurance, validation, and verification as constitutive features of any legitimate twin

Two research directions deserve explicit attention. The governance failure argument advanced here depends on an empirical claim, that organisations in this sector are systematically non-

Kustelega, Mekovec and Shareef (n 3) 1792; Ahmed and Uddin (n 1) 25; Parikh et al. (n 1) 23; Lennon, Julien and Quin (n 12) 1623-1624; Aloisi, Joppe and Abraha (n 6).

Lennon, Julien and Quin (n 12) 1619; Abraha (n 5) 40-42, on structural non-compliance persisting across jurisdictions primarily through exploitation of interpretive ambiguity and technical opacity, not through gaps in the law itself.

Teller (n 16) 1.

Helbing and Argota Sanchez-Vaquerizo (n 26) 9-10, 15.

Wagg et al. (n 36) 14-15

compliant, that is currently impossible to verify with precision. Vendor contract terms and DPIA records in this space are almost entirely unavailable for external analysis; what limited evidence exists comes from regulatory enforcement, which is itself sparse. Without that empirical foundation, the doctrinal argument is sound but its practical scale remains uncertain. The second direction is more immediately tractable. The Annex III classification of workforce digital twins under the EU AI Act currently rests on provider self-assessment, and the Commission has issued no guidance specifically addressed to HR analytics platforms. The compliance deadline deferral in the May 2026 Digital Omnibus creates a window before December 2027 enforcement in which that guidance gap could be addressed. Whether the Commission uses that window is an open question, and there is a reasonable case that scholarly analysis of this deployment category represents a practically useful contribution to the regulatory determination.

CONFLICTS OF INTEREST

The author declares no conflict of interest.

FUNDING ACKNOWLEDGEMENT

This research received no specific grant from any funding agency in the public, commercial, or not-

for-profit sectors.

REFERENCES

Primary Legal Sources

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) [2016] OJ L119/1.

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L2024/1689.

Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products [2024] OJ L2024/2853.

California Consumer Privacy Act of 2018 (Cal Civ Code ss 1798.100 et seq) as amended by the California Privacy Rights Act of 2020.

Illinois Biometric Information Privacy Act, 740 ILCS 14/1 et seq (2008).

ISO/IEC 30173:2023 Digital twin: Concepts and terminology (International Organisation for Standardization 2023).

Case C-634/21, SCHUFA Holding and Others (Scoring), Judgment of the Court of Justice of the European Union of 7 December 2023, ECLI:EU:C:2023:957.

Case C-203/22, D&B Austria GmbH v CK, Judgment of the Court of Justice of the European Union of 27 February 2025, ECLI:EU:C:2025:117.

Cothron v White Castle System Inc, 2023 IL 128004 (Supreme Court of Illinois, 2023).

Article 29 Working Party, "Guidelines on Automated Individual Decision-Making and Profiling for the Purposes of Regulation 2016/679" (WP 251 rev.01, 6 February 2018).

Secondary Sources

Abraha H, "Navigating Workers' Data Rights in the Digital Age: A Historical, Current, and Future Perspective on Workers' Data Protection" ILO Working Paper 149 (International Labour Office, Geneva, September 2025).

Ahmed MD and Uddin MJ, "Digital Twins and Legal Liability: Navigating Accountability in Simulated Realities" (2025) 7(8) International Journal of Law and Politics Studies 19.

Aloisi A, Joppe A and Abraha H, "From 'General' to 'Context-Specific' Data Protection for Workers: Insights from the EU Rules on Platforms, Algorithmic Management and AI Systems" (first published online 12 September 2025) European Labour Law Journal, DOI:10.1177/20319525251375031.

Bolyn A and Bechet E, "Identifying a Problem and Solving It with a Digital Twin Idea" (Digital Twin

- Academy Module, Interreg V-A Euregio Meuse-Rhine, April 2023).
- Brand D and Devlin S, "Ethical, Legal, and Technical Considerations for the Development of Digital Twins in Healthcare" (IEEE Conference Publication, August 2025)
DOI:10.1109/ICTCS63349.2025.11212607.
- BUSE Attorneys at Law, "Digital Twins of Real People: Legal Risks for Businesses" (BUSE Law Blog, March 2026) <<https://buse.de/en/blog-en/labor-law/data-protection/digital-twins-of-real-people-legal-risks-for-businesses/>>.
- Gore SI and others, "Ethical and Legal Implications of Digital Twin Deployment in Biomedical and Pharmaceutical Domains" in *Accelerating Product Development Cycles with Digital Twins and IoT Integration* (IGI Global Scientific Publishing 2025) 491.
- Helbing D and Argota Sanchez-Vaquerizo J, "Digital Twins: Potentials, Ethical Issues, and Limitations" [revised final version] (ETH Zurich / Complexity Science Hub Vienna 2022).
- Jørgensen BN and Ma ZG, "Digital Twins Under EU Law: A Unified Compliance Framework Across Smart Cities, Industry, Transportation, and Energy Systems" (2025) 14(24) *Electronics* 4881, DOI:10.3390/electronics14244881.
- Keerthana S, "Legal Frontiers of Digital Twins: Privacy, Ownership, and Responsibility" (2025) 5(1) *Journal of Legal Research and Juridical Sciences* 652.
- Kustelega M, Mekovec R and Shareef A, "Privacy and Security Challenges of the Digital Twin: Systematic Literature Review" (2024) 30(13) *Journal of Universal Computer Science* 1782.
- Lennon Y, Julien N and Quin A, "Ensuring Personal Data Compliance by Integrating Legal Constraints into Digital Twin Design Methodology" in *Proceedings of the 35th ESREL & 33rd SRA-Europe Conference* (Research Publishing, Singapore 2025) 1618.
- Mocanu D and Sibony AL, "EU Consumer Law Meets Digital Twins" (2023) 1 *Revue europeenne de droit de la consommation* 229.
- Nechesov A, Dorokhov I and Ruponen J, "Virtual Cities: From Digital Twins to Autonomous AI Societies" (2025) 13 *IEEE Access* 13866.
- Okan I, "AI Gets Personal: CCPA vs GDPR on Automated Decision-Making" (Berkeley Technology Law Journal, April 2025) <<https://btlj.org/2025/04/ccpa-vs-gdpr-on-automated-decision-making/>>.
- Parikh M, Kamath A, Jajesh V and Kapoor P, "The Tour d'Horizon of Data Law Implications of Digital Twins" (2025) 27 *OMG Journal of Innovation* (Nishith Desai Associates).
- Peldon D and others, "Navigating Urban Complexity: The Transformative Role of Digital Twins in Smart City Development" (2024) 111 *Sustainable Cities and Society* 105583.
- Teller M, "Legal Aspects Related to Digital Twin" (2021) 379(2207) *Philosophical Transactions of the Royal Society A* 20210023.

- Wagg DJ, Burr C, Shepherd J, Xuereb Conti Z, Enzer M and Niederer S, "The Philosophical Foundations of Digital Twinning" (2025) 6 Data-Centric Engineering e12.
- Zoltick MM and Maisel JB, "Societal Impacts: Legal, Regulatory and Ethical Considerations for the Digital Twin" in The Digital Twin (Springer International Publishing 2023) 1167.
- Rosin A and Parviainen H, "Improving the Transparency of Algorithmic Management in Platform Work" (2025) European Labour Law Journal DOI:10.1177/20319525251375023.
- Minotakis A, "Regulating AI in the Workplace" (2025) New Technology, Work and Employment DOI:10.1177/29768624251396248.
- EU Digital Omnibus on AI Provisional Agreement (7 May 2026), amending Regulation (EU) 2024/1689.
- Crowell & Moring LLP, "AI and HR in the EU: A 2026 Legal Overview" (24 February 2026).