EGE SCHOLAR JOURNAL

Ege Scholar Journal, Year: 2025, Volume: 2, Issue: 3, Pages : 184-195 **Research Article**

DOI: 10.5281/zenodo.18525816

Publication Date: 30.09.2025

Kostiantyn Lutsenko[1]

1. ORCID: https://orcid.org/0009-0007-4076-3028 Email: klutsenko@gmail.com Interregional Academy of Personnel Management, Ukraine Frometivska str., 2, Kyiv, 03039, Ukraine

Compliance Engineering for High-Risk HR AI under the EU AI Act: Discrimination Risks, Controls, and Audit Evidence

Abstract

Artificial intelligence (AI) is increasingly used in human resources for recruitment and worker evaluation, including résumé screening, candidate ranking, online assessments, video-interview scoring, and performance analytics. While these systems can improve efficiency and consistency, they may also introduce or amplify discrimination through proxy variables, historically biased labels, measurement error in "soft" constructs, and feedback loops across hiring and performance pipelines. This paper proposes a compliance engineering framework that operationalizes the EU Artificial Intelligence Act (Regulation (EU) 2024/1689) for high-risk HR AI systems by translating legal obligations into implementable technical and governance controls. The framework integrates the NIST AI Risk Management Framework lifecycle with HR-specific fairness practices, data protection safeguards relevant to automated decision-making, and enforceable bias-audit patterns from employment regulation. Results include (i) a reference governance-and-technical architecture for HR AI, and (ii) a control–metric matrix mapping discrimination risk modes to test procedures, mitigations, and audit-ready evidence artifacts. The paper concludes with practical compliance dossier templates suitable for both deployers and vendors, supporting traceability, meaningful human oversight, and continuous monitoring of performance and subgroup fairness.

Keywords: Algorithmic hiring; HR analytics; Algorithmic discrimination; EU AI Act; Compliance engineering; bias audit; Human oversight.

How to cite: Lutsenko K. (2025). Compliance Engineering for High-Risk HR AI under the EU AI Act: Discrimination Risks, Controls, and Audit Evidence. *Ege Scholar Journal, 2*(3), 184-195

Introduction

AI-mediated decision-making in human resources (HR) is no longer confined to back-office analytics; it increasingly shapes access to employment opportunities and career progression through automated résumé parsing and ranking, candidate sourcing and targeted advertising, online assessment scoring, video-interview analytics, and employee performance monitoring. These systems are adopted to reduce time-to-hire, improve consistency, and scale decision processes. However, HR decisions are high-stakes: errors can directly affect individuals' livelihoods and dignity, and discriminatory outcomes create material legal and reputational exposure.A central concern is algorithmic discrimination—systematic differences in outcomes for protected groups that are not justified by job-related necessity. A robust body of scholarship shows that data-driven systems can reproduce historical inequities embedded in training data, encode social bias through measurement and target-variable choices, and generate disparate impacts even when protected attributes are not explicitly included (Barocas & Selbst, 2016). In practice, HR data and feature pipelines often contain proxy variables (e.g., residential location, educational history, employment gaps) that correlate with protected characteristics; models can learn these correlations and reproduce differential outcomes through seemingly "neutral" predictors.In HR settings, discrimination risks are amplified by three recurring mechanisms. First, measurement error is common because many HR constructs (e.g., leadership potential, teamwork, "culture fit") are operationalized via noisy or weakly validated instruments, increasing the likelihood of construct invalidity and subgroup bias. Second, feedback loops arise when model-assisted decisions influence who is hired, trained, evaluated, and promoted, thereby shaping future training data and potentially locking in disparities unless monitored and corrected. Third, institutional incentives and vendor marketing can outpace rigorous validation: empirical analyses of algorithmic hiring vendors identify substantial variation in transparency, documentation of target variables, and the credibility of bias-mitigation claims (Raghavan et al., 2020).Against this backdrop, governance requirements for HR AI are tightening across jurisdictions. In the EU, the Artificial Intelligence Act establishes a risk-based regime with explicit obligations for high-risk AI systems, including those used in employment-related decision-making, covering risk management, data governance, technical documentation, logging, transparency, human oversight, and requirements for accuracy, robustness, and cybersecurity (European Union, 2024). In parallel, data protection law interacts strongly with HR AI: GDPR safeguards regarding automated decision-making and profiling, as reflected in EDPB-endorsed guidance, elevate expectations for transparency and meaningful safeguards where decisions produce legal or similarly significant effects (European Data Protection Board, n.d.). Outside the EU, enforcement signals similarly point toward auditable fairness controls. The U.S. Equal Employment Opportunity Commission emphasizes that existing anti-discrimination laws apply when AI is used in employment decisions (EEOC, 2024), while New York City's Local Law 144 requires independent bias audits and public disclosures for certain automated employment decision tools (New York City DCWP, n.d.).This paper treats the EU AI Act as a benchmark for compliance engineering: translating legal obligations into concrete technical controls, evidence artifacts, and measurable metrics implementable by both employers (deployers) and vendors (providers). We integrate the NIST AI Risk Management Framework (AI RMF 1.0) as a lifecycle structure for governing, mapping, measuring, and managing risks (NIST, 2023), and incorporate ISO/IEC TR 24027 guidance on identifying and addressing bias vulnerabilities across AI lifecycle phases (ISO/IEC, 2021). The result is an implementable governance-and-technical architecture and a control–metric matrix that connects discrimination failure modes to testable mitigations and audit-ready evidence.

Research objectives

1. **To characterize discrimination and error modes** in AI-based recruitment and worker evaluation pipelines, and link these modes to **measurable harms** and fairness outcomes. (Barocas & Selbst, 2016)
2. **To propose an end-to-end reference architecture** for HR AI governance and technical controls aligned with **high-risk system obligations** under the EU Artificial Intelligence Act (Figure 1). (European Union, 2024)
3. **To specify a control–metric matrix** mapping discrimination risk modes to **test procedures, mitigations, and audit-ready evidence artifacts** suitable for employers and vendors (Table 1). (NIST, 2023; New York City DCWP, n.d.)
4. **To provide implementable compliance artifacts** (documentation templates, logging specifications, and human-oversight procedures) consistent with the **NIST AI RMF lifecycle** and relevant **GDPR safeguards** for automated decision-making. (European Data Protection Board, n.d.; NIST, 2023)

Materials and Methods

Materials: Authoritative legal texts, standards, and research sources

This study draws on authoritative legal instruments, standards, and peer-reviewed or widely cited policy research to construct a compliance engineering framework for HR AI.

EU AI Act benchmark.

The EU Artificial Intelligence Act (Regulation (EU) 2024/1689) was used as the primary regulatory benchmark, including its provisions on high-risk systems and the employment-related high-risk use cases listed in Annex III (European Union, 2024). Where interpretive context was needed, the analysis prioritized official legal text and reputable legal commentary that maps employment AI use cases to high-risk obligations.

Risk management and bias standards.

To operationalize legal requirements into implementable controls, the framework incorporates the NIST AI Risk Management Framework (AI RMF 1.0) as a lifecycle-oriented risk management structure (NIST, 2023) and ISO/IEC TR 24027:2021 as a reference for identifying and addressing bias-related vulnerabilities across the AI lifecycle (ISO/IEC, 2021).

Data protection and automated decision-making.

Because HR AI frequently intersects with automated decision-making safeguards, GDPR provisions relevant to automated decision-making were treated as complementary constraints, along with EDPB-endorsed guidance on automated decision-making and profiling (European Data Protection Board, n.d.; European Union, 2016).

External enforcement and governance signals.

To inform implementability and audit expectations, the study considered enforcement signals and operational models from jurisdictions where HR AI governance is actively developing, including EEOC technical assistance emphasizing that existing anti-discrimination laws apply to AI in employment (EEOC, 2024) and the audit and notice regime in New York City's Local Law 144 for automated employment decision tools (New York City DCWP, n.d.).

Research foundations on algorithmic discrimination and accountability.

The conceptualization of discrimination pathways and auditability builds on foundational scholarship on disparate impact in data-driven systems (Barocas & Selbst, 2016), empirical evidence on vendor practices and bias-mitigation claims in algorithmic hiring (Raghavan et al., 2020), policy analysis of hiring algorithms and equity (Upturn, 2018), and accountability-by-design principles for algorithmic systems (Kroll et al., 2017).

Methods: Compliance engineering procedure

The method follows a structured compliance engineering approach that translates legal and standards-based requirements into testable controls, measurable metrics, and audit-ready evidence artifacts.

Step 1: HR decision decomposition.

HR AI use was decomposed into decision stages—sourcing, screening, assessment, interview evaluation, offer/selection, and performance evaluation. Each stage was analyzed for discrimination vectors (proxy discrimination, measurement error, feedback loops, and process-level harms) and for the governance requirements that determine how model outputs translate into employment actions.

Step 2: Risk and obligation mapping.

High-risk obligations under the EU AI Act were mapped into implementable control families, including: (i) risk management and residual risk acceptance, (ii) data governance and data quality, (iii) technical documentation and record-keeping/logging, (iv) transparency and user information duties, (v) human oversight, and (vi) accuracy, robustness, and cybersecurity controls (European Union, 2024). This mapping produced an obligation-to-control traceability structure suitable for procurement, deployment, and audit.

Step 3: Lifecycle alignment.

Mapped controls were aligned to NIST AI RMF functions (Govern, Map, Measure, Manage) to ensure lifecycle coverage and continuous monitoring rather than one-time compliance activities (NIST, 2023). This alignment supports iterative improvement through periodic evaluation, incident response, and controlled change management.

Step 4: Bias audit operationalization.

Operational bias audit expectations from New York City Local Law 144 were incorporated as an illustrative enforceable model for periodic auditing, documentation, and notice practices (New York City DCWP, n.d.). The intent is not to transpose LL144 requirements universally, but to use it as a practical template for translating governance expectations into recurring audit cycles and public/internal evidence.

Step 5: Output artifacts.

Two implementable artifacts were produced: (i) a reference governance-and-technical architecture for HR AI compliance engineering (Figure 1), and (ii) a control–metric matrix linking discrimination risk modes to controls, evidence artifacts, and measurable indicators (Table 1). Together, these outputs support audit-readiness and cross-stakeholder accountability for providers and deployers.

Results

Result 1. Employment AI is compliance-sensitive by design (not by add-on)

Because HR AI mediates access to employment opportunities and employment terms, compliance cannot be treated as a late-stage overlay. The EU AI Act's high-risk regime requires that HR AI systems be designed for traceability, documentation, logging, transparency, meaningful human oversight, and robustness across their lifecycle. In parallel, GDPR safeguards for automated decision-making reinforce expectations for contestability and procedural safeguards where decisions have significant effects. (European Union, 2024; European Data Protection Board, n.d.).

Figure 1 (mandatory)

Figure1.ComplianceengineeringarchitectureforHRAI(EUAIActbenchmark+NISTAIRMFlifecycle) This reference architecture operationalizes high-risk themes under the EU AI Act and aligns them to the NIST AI RMF lifecycle functions (GOVERN–MAP–MEASURE–MANAGE), enabling continuous monitoring, auditable evidence, and controlled improvement. (European Union, 2024; NIST, 2023).

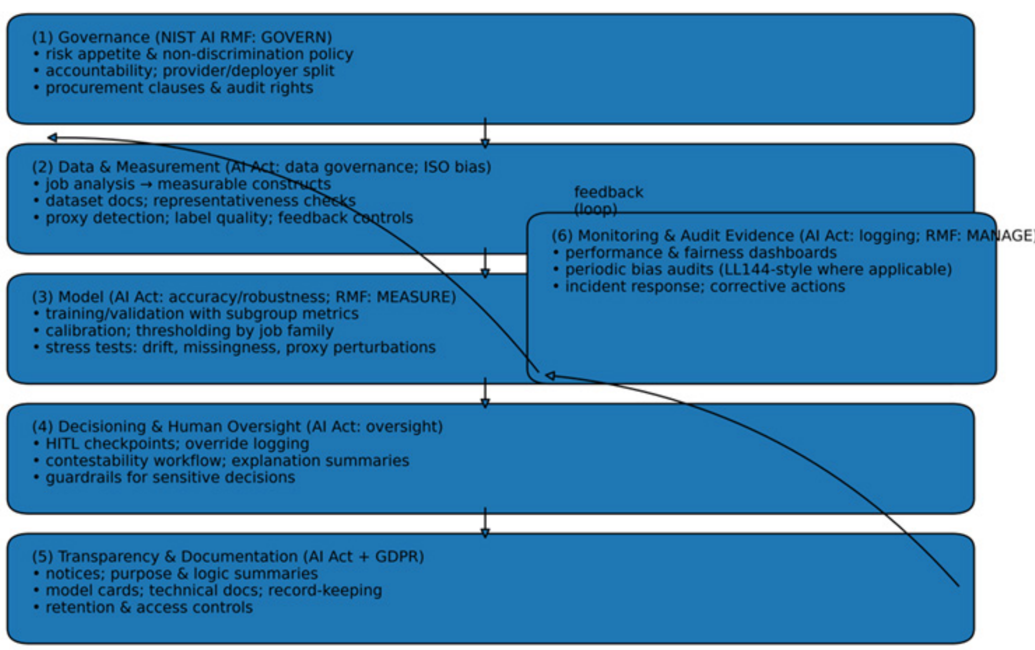

In-text callout example: "The end-to-end compliance engineering architecture is shown in Figure 1."

Table 1 (mandatory)

Table 1. Discrimination and compliance risk modes in HR AI, controls, and measurable metrics

Risk mode	Typical HR AI context	Mechanism	Controls (technical + governance)	Audit evidence artifacts	Metrics (examples)
Proxy discrimination	Sourcing, résumé screening, ranking	Inputs correlate with protected traits; model reproduces group differences via proxies	Proxy testing and sensitivity analysis; feature register; feature review board; prohibition list for high-risk features	Feature register; proxy test report; change-control approvals	Selection-rate ratios; disparate impact indicators; proxy sensitivity deltas
Historical bias in labels	Performance prediction; "high potential" classification	Biased human ratings become "ground truth" labels	Label governance; rater calibration; audit sampling; label quality checks	Label SOP; inter-rater reliability report; label audit log	Inter-rater agreement; subgroup label shift; error rate by subgroup
Measurement error / construct invalidity	"Culture fit," soft-skill scoring, video-interview analytics	Noisy instruments; weak construct validity; subgroup measurement bias	Job analysis; validated assessments; construct documentation; periodic validity review	Job analysis report; validity study; model card sections on construct limits	Predictive validity by subgroup; calibration by subgroup; false positive/negative rates
Feedback loops	Hiring → performance → retraining pipeline	Model influences workforce composition and future labels	Cohort tracking; delayed-outcome monitoring; periodic re-baselining; drift triggers	Cohort monitoring report; model change logs; retraining triggers	Drift in subgroup mix; outcome stability; longitudinal fairness trends
Over-automation / de facto sole automation	Adverse screening/ selection decisions	Human review is nominal; decisions become effectively automated	Mandatory human oversight gates; override capability; reviewer training; "do not automate" rules	Oversight policy; override logs; reviewer training records	% adverse decisions reviewed; override rate; time-to-review
Opacity / contestability failure	Candidate rejection or adverse employment action	Inadequate notice/ explanation; weak appeal pathway	Structured notices; explanation summaries; appeal workflow; documentation retention	Notice templates; appeals log; decision trace package	Appeal turnaround time; explanation completeness; reversal rate
Vendor accountability gap	Procured HR AI tools	Unclear provider/ deployer responsibilities; limited auditability	Contract clauses defining obligations; audit rights; documentation deliverables; incident cooperation	Procurement clauses; assurance reports; vendor disclosure pack	Vendor disclosure score; audit pass rate; incident response SLAs
Bias-audit non-compliance	Regulated jurisdictions (e.g., NYC)	Missing audit, disclosure, or notice requirements	Annual independent bias audit; public summary; candidate/employee notices	Bias audit report; public posting link; notice logs	Audit recency; notice compliance rate; remediation closure rate

Interpretation. Table 1 translates common discrimination and governance failure modes into testable controls and audit-ready evidence. This reflects the practical direction of current governance frameworks: lifecycle risk management (NIST AI RMF) and enforceable audit/notice patterns (e.g., NYC AEDT regime). (NIST, 2023; New York City DCWP, n.d.; Raghavan et al., 2020).

3.1 Algorithmic discrimination mechanisms in recruitment and evaluation (revised)

Discrimination pathways in HR AI are typically structural rather than intentional. First, proxy discrimination arises when input features correlate with protected traits; even if protected attributes are excluded, models can learn correlational patterns that yield disparate impacts (Barocas & Selbst, 2016). Second, measurement and construct-validity issues are pervasive because HR constructs such as "leadership potential" or "team fit" are difficult to operationalize; when these constructs are approximated through noisy behavioral signals or historically biased evaluations, subgroup measurement bias becomes likely. Third, vendor practice variability can weaken safeguards: empirical research documents inconsistent transparency and uneven validation of bias-mitigation claims across algorithmic hiring vendors (Raghavan et al., 2020). Fourth, feedback loops can lock in inequality: hiring decisions shape future workforce composition and performance labels, shifting the training distribution and potentially reinforcing disparities unless monitored and corrected. Fifth, process-level harms can occur even when predictive accuracy is acceptable, such as disproportionate manual-review "friction" imposed on specific groups. Consequently, fairness assessment should cover the full workflow—access to opportunities, selection rates, time-to-decision, review rates, and downstream outcomes—rather than only model-level accuracy.These mechanisms imply that technical "debiasing" alone is insufficient. Effective practice requires governance: explicit role accountability, documented job analysis, disciplined feature governance, and continuous auditing supported by evidence artifacts (Kroll et al., 2017).

3.1.1 EU AI Act benchmark: translating high-risk obligations into HR controls (revised)

The EU AI Act establishes a risk-based framework and identifies "high-risk" AI systems through Article 6 criteria and Annex III use cases, which include employment-related contexts such as recruitment and worker management. For such systems, obligations on risk management, data governance, documentation/logging, transparency, human oversight, and robustness/cybersecurity can be operationalized into concrete HR controls and audit evidence. (European Union, 2024).

A compliance engineering interpretation is as follows:

1. **Risk management:** Define acceptable error and fairness thresholds by job family, document residual risks, and maintain escalation and remediation pathways for discrimination incidents. (European Union, 2024).

2. **Data governance:** Maintain dataset documentation, representativeness checks, label-quality governance, and bias assessments aligned to lifecycle bias guidance. (ISO/IEC, 2021).

3. **Documentation and logging:** Produce model cards/technical documentation and event logs that support post-hoc review of adverse outcomes and change control. (European Union, 2024).

4. **Transparency and human oversight:** Ensure meaningful review points and avoid de facto solely automated adverse decisions where significant effects arise; implement safeguards consistent with ADM/profiling guidance. (European Data Protection Board, n.d.).

5. **Robustness and cybersecurity:** Protect scoring pipelines from tampering, monitor drift and missingness, and maintain controlled release processes for model updates. (European Union, 2024).

In practice, this means HR AI governance should be engineered like a regulated product: controlled change management, audit-ready evidence, and documented human oversight.

Discussion

Compliance engineering as a socio-technical discipline

The results support a core proposition: fairness and compliance in HR AI are socio-technical properties. Discrimination risk arises not only from model parameters but also from organizational processes, including how jobs are defined, how "success" is measured, how performance evaluations are conducted, how features are selected, and how reviewers interpret model outputs. Accordingly, effective governance must address the full decision pipeline—not merely the model artifact—because upstream measurement choices and downstream human actions can dominate observed disparities.The EU AI Act provides a strong benchmark precisely because it forces organizations to connect governance commitments to verifiable technical evidence. By treating employment-related AI as potentially high-risk, the Act implicitly recognizes that HR AI can affect fundamental rights and life opportunities at scale (European Union, 2024). When organizations adopt the Act as a "design target," they can standardize auditable controls beyond the EU market and improve cross-jurisdiction readiness, particularly when procurement and vendor management are structured around evidence deliverables and change-control obligations.

Interaction with GDPR and automated decision-making safeguards

Data protection safeguards are central to HR AI governance because adverse employment outcomes can constitute "similarly significant effects." The GDPR framework for automated decision-making and profiling, together with EDPB-endorsed guidance, elevates expectations for transparency and meaningful safeguards, including contestability and human intervention where applicable (European Data Protection Board, n.d.; European Union, 2016). A practical implication is that "human-in-the-loop" must be operationally meaningful—review should be competent, documented, and capable of changing outcomes. This requires reviewer training, clear escalation procedures, and evidence that review is not a rubber stamp (e.g., override logs, reversal rates, and documented rationales).

External enforcement signals and operational bias audits

Regulatory and enforcement signals outside the EU indicate convergent expectations: organizations will be asked to demonstrate that fairness controls are real, measured, and auditable. In the United States, the EEOC reiterates that existing anti-discrimination laws apply when AI is used in employment decisions (EEOC, 2024). At the municipal level, New YorkCity's Local Law 144 operationalizes compliance through periodic bias audits and notice/disclosure requirements for certain automated employment decision tools (New York City DCWP, n.d.). While LL144 differs in scope and legal foundation from the EU AI Act, it illustrates a governance direction likely to diffuse: recurring independent checks, documented mitigations, and user-facing transparency.

Limits of purely technical mitigation

Technical bias mitigation (e.g., reweighing, constrained optimization, post-processing) can reduce metric disparities under defined assumptions, but it does not address construct invalidity, biased labels, or proxy pathways embedded in upstream HR processes. Empirical research on algorithmic hiring vendors underscores that bias-reduction claims can vary widely in evidentiary quality and transparency (Raghavan et al., 2020). Consequently, compliance engineering prioritizes (i) job-related construct definition, (ii) instrument validity and label governance, (iii) proxy and feature governance with documented decision rights, and (iv) continuous monitoring supported by auditable evidence artifacts. ISO/IEC TR 24027 supports this lifecycle view by emphasizing bias vulnerabilities and mitigations across the system lifecycle rather than only at model-training time (ISO/IEC, 2021).

Limitations and future work

This study is a design-oriented synthesis rather than a causal empirical evaluation. The proposed controls and artifacts should be validated through deployments that measure both utility (e.g., time-to-hire, predictive validity) and fairness outcomes (e.g., subgroup selection rates, error rates, calibration) pre- and post-mitigation, with careful attention to sample sizes, uncertainty intervals, and shifting labor-market distributions. Future work should also assess how different HR decision stages (sourcing vs. screening vs. interview analytics) require distinct evidentiary thresholds and oversight intensity, particularly for high-impact adverse decisions.

Conclusions

AI can improve efficiency and consistency in HR processes, but it also concentrates discrimination risk into scalable decision pipelines. The central conclusion of this paper is that fairness and compliance cannot be treated as post-hoc model tuning; they require end-to-end compliance engineering aligned to high-accountability governance expectations.Using the EU AI Act as a benchmark, HR AI used in recruitment and evaluation should be engineered as a high-risk, audit-ready system with traceability by design. The Act's obligations—risk management, data governance, technical documentation, logging, transparency, human oversight, and requirements for accuracy, robustness, and cybersecurity—translate into implementable controls: job analysis, dataset documentation, fairness evaluation by subgroup, proxy and feature governance, meaningful human review, and evidence retention (European Union, 2024). The NIST AI RMF provides a lifecycle structure for governing, mapping, measuring, and managing these risks over time, including incident response and controlled model change (NIST, 2023). ISO/IEC TR 24027 strengthens this approach by focusing on bias vulnerabilities across lifecycle phases and highlighting the importance of measurement choices and process controls (ISO/IEC, 2021). GDPR safeguards and EDPB guidance further reinforce the need for transparency and meaningful human involvement where automated decision-making significantly affects individuals (European Data Protection Board, n.d.; European Union, 2016).Operationally, organizations should maintain a compliance-ready HR AI dossier, including a feature register, subgroup fairness evaluation reports, oversight logs, contestability workflows, and vendor contract clauses with audit rights and documentation deliverables. External signals—including EEOC guidance and the operational bias-audit model introduced by NYC Local Law 144—indicate increasing demand for demonstrable, auditable fairness practices rather than aspirational policies (EEOC, 2024; New York City DCWP, n.d.). An evidence-first approach—implemented via the architecture in Figure 1 and the control–metric matrix in Table 1—supports HR AI systems that are not only operationally effective but also legally defensible and aligned with emerging international norms for trustworthy AI in employment.

Acknowledgments

The author acknowledges the contributions of public institutions and standards bodies that provide foundational guidance for trustworthy AI and fair employment decision-making, including the European Union (EU AI Act), NIST (AI RMF), the European Data Protection Board (guidance on automated decision-making), ISO/IEC (bias guidance), and public regulators developing governance requirements for automated employment decision tools (European Union, 2024; NIST, 2023; European Data Protection Board, n.d.; ISO/IEC, 2021; New York City DCWP, n.d.).

Appendix A. Minimum HR AI compliance dossier (audit-ready artifacts)

A1. Job analysis and construct definition: Define what the model predicts, why it is job-related, and known limitations.

A2. Dataset documentation: Sources, representativeness, label quality, retention and access controls (ISO/IEC, 2021).

A3. Feature register and proxy governance: Proxy-risk assessment, restricted features list, and review board sign-off (Barocas & Selbst, 2016).

A4. Fairness evaluation report: Subgroup selection rates, error rates, calibration, and "friction" metrics across the workflow.

A5. Human oversight SOP: Review triggers, override rules, training requirements, logging, and escalation pathways (European Data Protection Board, n.d.).

A6. Transparency package: Candidate notices, explanation summaries, and appeal/contestability workflow with retention (European Data Protection Board, n.d.).

A7. Monitoring and change management plan: Periodic audit cadence, drift triggers, incident response, and controlled releases (NIST, 2023).

Appendix B. Bias-audit protocol (LL144-style operationalization + EU AI Act benchmark)

B1. Define the decision points in scope (screening, ranking, assessment scoring, interview analytics) (New York City DCWP, n.d.).

B2. Define legally applicable protected-group categories for analysis; compute selection rates and disparity ratios.

B3. Compute uncertainty estimates (e.g., confidence intervals) and sample-size checks; document data limitations.

B4. Perform proxy sensitivity tests (remove/perturb suspect features; quantify outcome shifts) (Barocas & Selbst, 2016).

B5. Document mitigation actions (threshold changes, feature removal, retraining, added human review gates) and verify post-mitigation performance.

B6. Publish required summaries where mandated; maintain internal evidence packs for audits and incident response (New York City DCWP, n.d.).

References

1. European Data Protection Board. (2018). *Automated decision-making and profiling*.
2. European Data Protection Board. (2020). *Guidelines 05/2020 on consent under Regulation (EU) 2016/679*.
3. European Parliament and Council of the European Union. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union*, L 119, 1–88.
4. European Union. (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act).

5. Equal Employment Opportunity Commission. (n.d.). *The EEOC's role in artificial intelligence*.

6. Future of Privacy Forum. (2022). *Automated decision-making under the GDPR: Practical cases from courts and data protection authorities*.

7. Hunton Andrews Kurth LLP. (n.d.). *Impact of the EU Artificial Intelligence Act on human resources activities*.

8. ISO/IEC. (2021). *ISO/IEC TR 24027: Information technology—Artificial intelligence—Bias in AI systems and AI-aided decision making*. International Organization for Standardization.

9. Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165(3), 633–705.

10. Mayer Brown. (n.d.). *Commentary on EEOC guidance concerning AI and disparate impact*.

11. New York City Department of Consumer and Worker Protection. (2021). *Local Law 144 of 2021: Automated Employment Decision Tools*.

12. New York City Department of Consumer and Worker Protection. (2023). *Rules implementing Automated Employment Decision Tools (Statement of basis and purpose)*.

13. New York State Office of the State Comptroller. (2025). *Enforcement of Local Law 144: Automated Employment Decision Tools* (Audit report).

14. Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. In *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency* (pp. 469–481).

15. Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2019). *Mitigating bias in algorithmic hiring: Evaluating claims and practices* (Working paper).

16. Reuters. (2024). EEOC litigation concerning alleged bias in AI-driven hiring software.

17. Reuters. (2024). Legal challenge to algorithmic hiring systems under U.S. employment discrimination law.

18. Taylor Wessing. (n.d.). *The European Union Artificial Intelligence Act from a human resources perspective*.

19. Upturn. (2018). *Help wanted: An exploration of hiring algorithms, equity, and bias*.

20. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671–732.

21. Article 29 Data Protection Working Party. (2018). *Guidelines on automated individual decision-making and profiling under Regulation (EU) 2016/679* (WP251rev.01).

22. European Data Protection Board. (n.d.). *Endorsed Article 29 Working Party guidance on automated decision-making*.

23. Organisation for Economic Co-operation and Development. (n.d.). *General guidance on trustworthy artificial intelligence*.

24. National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*.

25. New York City Department of Consumer and Worker Protection. (n.d.). *Automated Employment Decision Tools: Regulatory guidance*.

26. Perkins Coie LLP. (2023). *Final rules under New York City Local Law 144 governing automated employment decision tools*.

27. International Association of Privacy Professionals. (2022). *Coverage and analysis of automated decision-making under the GDPR*.

28. European Union. (2016). Regulation (EU) 2016/679, Article 22: Automated individual decision-making, including profiling.

29. ISO/IEC. (2021). *Artificial intelligence—Bias in AI systems and AI-aided decision making* (Technical report).

30. Future of Privacy Forum. (2022). *Automated decision-making under the GDPR*.