



Governing Fair Algorithmic Hiring Under the EU AI Act: A Multi-Layer Decision Architecture for High-Risk Recruitment Systems

Veranika Khlud^{1,2,3*}

¹ Riga Technical University, Institute of Mechanical Engineering, Aerospace Technologies and Transport, Riga, Latvia

² Baltic International Academy, Riga, Latvia

³ International Transport Development Association, Riga, Latvia

ARTICLE INFO

Article history:

Received 29 December 2025

Received in revised form 5 February 2026

Accepted 15 March 2026

Available online 21 April 2026

Keywords:

Algorithmic hiring, EU AI Act, AI governance, Organizational justice, Algorithmic fairness, Human oversight

ABSTRACT

Artificial intelligence is increasingly used to support recruitment decisions, while the European Union's (EU) Artificial Intelligence (AI) Act treats many employment-related AI systems as high-risk because they may significantly affect career opportunities, livelihoods, and fundamental rights. However, existing debates on fair algorithmic hiring often remain fragmented, treating fairness either as a technical bias-mitigation issue, a legal compliance issue, or a matter of applicant perception. This fragmentation leaves organizations without an integrated governance architecture that connects regulatory obligations, decision accountability, human oversight, data controls, and procedural justice. The purpose of this conceptual paper is to develop a multi-layer decision architecture for governing fair algorithmic hiring under the EU AI Act. The article applies a theory-building methodology based on conceptual synthesis, integrating the AI Act's requirements for high-risk systems with organizational justice theory, algorithmic fairness research, and human resource management governance literature. The study develops a six-layer governance architecture consisting of strategic accountability, lifecycle risk management, data and model assurance, human oversight, candidate-facing transparency and redress, and continuous monitoring and organizational learning. In addition, the article formulates seven theoretical propositions explaining how governance mechanisms can translate compliance obligations into distributive, procedural, interpersonal, and informational justice in recruitment. The article concludes that fair algorithmic hiring cannot be achieved through isolated audits, bias testing, or disclosure statements alone. Instead, it requires an integrated decision architecture that allocates rights and responsibilities across the recruitment pipeline, preserves meaningful human judgment, enables candidate voice and appeal, and embeds legal and ethical safeguards into ongoing organizational routines.

1. Introduction

Algorithmic hiring has moved from experimental human resource (HR) technology to a core component of contemporary recruitment. Applicant tracking systems, automated ranking tools,

* Corresponding author.

E-mail address: veranika.khlud@itda.co

<https://doi.org/10.66972/ada21202624>

© The Author(s) 2026 | [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/)

chatbots, structured video interviews, psychometric games and machine-learning systems are used to source, screen, assess and prioritize candidates. These tools promise speed, scale and consistency, but they also alter the structure of employment decisions. A recruitment decision is no longer made only by a recruiter applying professional judgment to candidate information. It is increasingly produced by a socio-technical decision pipeline in which vendors, employers, data scientists, HR managers, legal departments and applicants all interact with a system whose outputs shape access to work [1-3].

The governance relevance of algorithmic hiring has intensified because the European Union's (EU) Artificial Intelligence (AI) Act classifies many AI systems used in employment, worker management, and access to self-employment as high-risk when they are used for recruitment, selection, decisions affecting work-related relationships, task allocation, monitoring or evaluation [4]. The regulation requires high-risk AI systems to satisfy obligations concerning risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness and cybersecurity [4]. The Act also introduces obligations for deployers, including appropriate use, competent human oversight, monitoring, log retention and information duties toward affected persons and workers' representatives [4]. At the time of writing, the implementation timetable remains progressive and subject to ongoing simplification debates, but the policy direction is clear: employment-related AI is no longer a voluntary ethics matter; it is a regulated decision domain [5].

Despite this regulatory shift, the academic and managerial literature remains fragmented. One stream examines the technical and legal limits of bias mitigation in algorithmic hiring, showing that vendor claims about objectivity often exceed what existing systems can demonstrate [2-3, 6-7]. A second stream examines human resource management (HRM) adoption, emphasizing efficiency, analytics capability and organizational transformation [1, 8-10]. A third stream focuses on applicant and employee reactions, showing that algorithmic evaluation may influence trust, dignity, privacy concerns and willingness to engage with an employer [11-14]. A fourth stream studies algorithmic fairness and accountability more broadly, including data documentation, model transparency, audits and the limits of abstract fairness metrics [15-19]. These streams provide important insights, but they do not yet offer a unified architecture for governing fair hiring decisions across the full recruitment lifecycle.

The main weakness of the current debate is that fairness is frequently treated as a narrow property of a model, a one-time audit result, or a disclosure obligation. This is insufficient for hiring. A technically debiased model may still be unfair if applicants cannot understand the process, contest errors, receive respectful treatment, or benefit from meaningful human review. Similarly, a legally compliant system may still be perceived as illegitimate if governance routines are opaque and candidate voice is absent. Organizational justice theory has long shown that fairness in selection depends not only on decision outcomes but also on the perceived quality of procedures, explanations, interpersonal treatment and opportunities for voice [20-23]. Algorithmic hiring therefore requires a governance model that connects technical accuracy and bias controls with procedural, interpersonal and informational justice.

This article addresses an important gap in the literature by developing a governance architecture for fair algorithmic hiring under the EU AI Act. Existing research provides valuable insights into bias mitigation, applicant reactions, compliance requirements and fairness metrics, but it does not yet explain how these elements should be integrated into a coherent organizational governance model across the recruitment lifecycle. The significance of this study is twofold. First, it translates the legal architecture of high-risk artificial intelligence into organizational decision rights, control points and feedback loops. Second, it extends algorithmic fairness research by embedding fairness metrics

within a broader HRM governance model grounded in organizational justice. The objective of this study is to develop a conceptual framework specifying how organizations should design governance architectures that convert AI Act compliance into fair and legitimate algorithmic hiring decisions. To achieve this objective, the article proposes a six-layer decision architecture and seven theoretical propositions that explain how regulatory, technical and organizational mechanisms interact to produce fairness in high-risk recruitment systems.

2. Methodology

This study adopts a conceptual and theory-building approach. The methodological aim is to develop a model that connects previously separate research domains and clarifies the mechanisms through which fair algorithmic hiring can be governed. Conceptual articles require explicit research design, transparent selection of source domains and a clear account of the intended contribution [24]. This article follows a model-building approach: it specifies the focal phenomenon, defines the core constructs, synthesizes compatible theoretical lenses and derives propositions that can guide future empirical work [24-26].

The focal phenomenon in this study is algorithmic hiring, defined as recruitment processes in which an artificial intelligence system materially influences candidate access to work by sourcing, screening, ranking, assessing, recommending, or otherwise informing recruitment decisions. This definition includes both fully automated systems and hybrid human–artificial intelligence systems. It excludes administrative tools that do not materially affect candidate evaluation, although cases at the boundary of this definition may still raise transparency or data protection concerns. In turn, governance architecture refers to an organized set of decision rights, accountability roles, documentation practices, technical controls, oversight routines, and feedback mechanisms that structure how algorithmic hiring systems are procured, configured, deployed, and monitored.

The conceptual synthesis proceeded in four steps. First, the article identified regulatory design requirements from the EU AI Act, especially the high-risk system requirements and deployer obligations relevant to employment decisions [4, 5]. Second, it selected organizational justice theory as the normative and behavioral lens for fairness in hiring, distinguishing distributive, procedural, interpersonal and informational justice [20-23]. Third, it integrated algorithmic fairness and accountability research, including work on disparate impact, fairness trade-offs, documentation, auditing and the limits of abstraction [15-19, 27-28]. Fourth, it incorporated HRM governance and AI-in-HR literature to locate controls within the recruitment pipeline and the organization rather than only within the technical model [1, 8-10, 29-31].

The integration logic is abductive. Rather than treating legal compliance, technical fairness, and applicant justice as separate objectives, the article conceptualizes them as mutually dependent components of the same decision system. The AI Act supplies minimum legal obligations and lifecycle requirements. Organizational justice theory explains why candidates and employees may accept or reject algorithmic hiring as fair. Algorithmic fairness research identifies measurement and design limits. HRM governance identifies the organizational actors, incentives, and routines through which the system is enacted in practice. The resulting architecture is therefore not a checklist but a decision model that links regulatory duties to fairness mechanisms across the recruitment lifecycle.

Because the article is conceptual, the term results refer to the conceptual outputs of the synthesis: a structured governance architecture, a mapping of regulatory requirements onto governance functions, and a set of theoretical propositions. These outputs are designed to support empirical scrutiny in subsequent research. Specifically, future studies may test whether candidate appeal mechanisms improve perceived procedural justice, whether human oversight is effective only

when reviewers possess sufficient authority and competence, and whether fairness deteriorates in the absence of post-deployment monitoring.

3. Results

Algorithmic hiring should be understood as a sequence of interdependent decisions rather than a single automated recommendation. A recruitment pipeline begins with job design and advertising, continues through sourcing, screening, assessment and ranking, and ends with selection, feedback and possible appeal. Each stage creates a distinct fairness risk. Job advertisements may reproduce occupational stereotypes; sourcing tools may exclude candidates whose profiles are less visible online; screening systems may rely on proxies for elite education or prior employment; assessments may underrepresent non-standard career paths; ranking tools may compress heterogeneous candidates into a single score; and automated rejections may deprive candidates of feedback or voice [2-3, 6, 29-31].

The EU AI Act provides a legal baseline for controlling some of these risks. Articles [8-12, 29-31] require high-risk systems to satisfy requirements concerning risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness and cybersecurity [4]. For deployers, Article 26 is especially important because it shifts governance from the provider's technical design to the employer's organizational use of the system. Employers deploying high-risk AI systems must use them according to instructions, assign human oversight to competent and authorized people, monitor operation, keep logs where under their control, inform affected workers in workplace contexts and inform natural persons when high-risk AI assists decisions related to them [4].

However, legal obligations do not automatically produce fair hiring. Fairness in recruitment involves at least four justice dimensions. Distributive justice concerns the fairness of outcomes, including whether candidates from different groups have equitable access to opportunities. Procedural justice concerns the fairness of the decision-making process, including consistency, job-relatedness, bias suppression, voice and correctability. Interpersonal justice concerns whether applicants are treated with respect and dignity. Informational justice concerns the quality and sincerity of explanations [20-23]. Algorithmic hiring can undermine any of these dimensions even when a model appears statistically accurate. For example, recent evidence shows that algorithmic evaluation may be perceived as lacking respect and individualized consideration, not only as biased [12].

The architecture proposed in this article is built on the premise that fair algorithmic hiring requires alignment between legal compliance, technical assurance and organizational justice. Figure 1 presents the overall model. The architecture contains six layers: strategic accountability, lifecycle risk management, data and model assurance, human oversight, candidate-facing transparency and redress, and continuous monitoring and learning. These layers operate across the recruitment pipeline and are linked by feedback loops that convert incidents, applicant responses and monitoring results into system improvement.

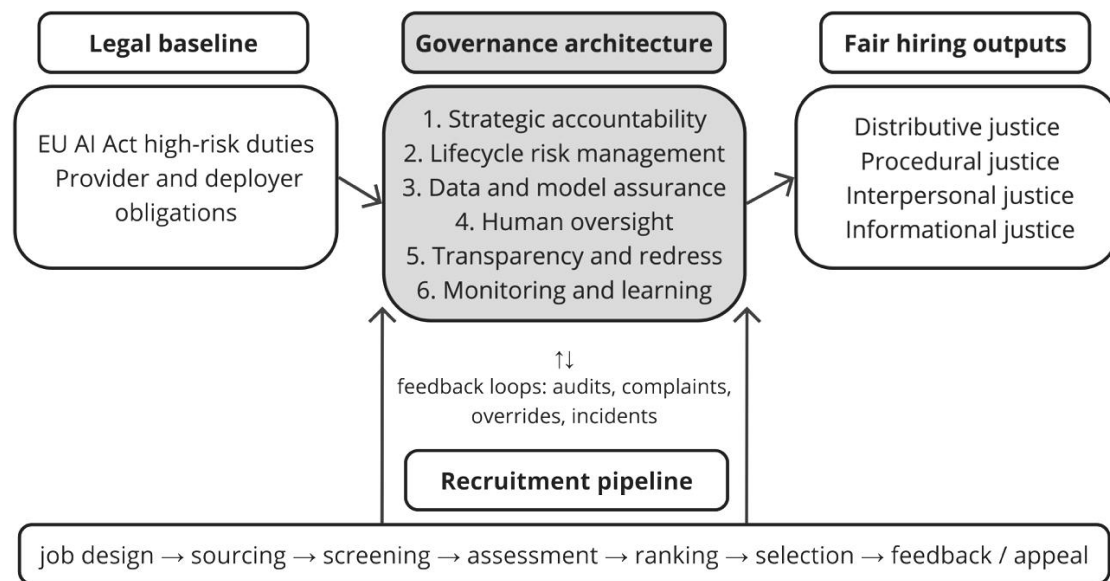


Fig. 1. Multi-layer governance architecture for fair algorithmic hiring under the EU AI Act

The first layer, strategic accountability, defines why the organization uses algorithmic hiring, which decisions may be delegated to AI, who is accountable for outcomes, and which fairness principles are non-negotiable. This layer prevents the common governance failure in which AI procurement is treated as an HR technology purchase rather than as a decision-system redesign. Strategic accountability should include a recruitment AI charter, explicit decision rights, escalation pathways, vendor accountability clauses and periodic board- or senior-management review. It also establishes the organization's interpretation of fair access to employment, including commitments to non-discrimination, job-relatedness, inclusion and candidate dignity.

The second layer, lifecycle risk management, translates the AI Act's risk-management logic into HR routines. It begins before procurement and continues after deployment. Before adoption, the employer should identify the intended purpose of the system, the recruitment stage affected, the candidate groups likely to be affected, the expected benefits, the risks of harm and the alternatives to automation. During procurement and configuration, the employer should require evidence of validation, data governance, model limitations, documentation, accuracy and cybersecurity from the provider. During use, the employer should monitor error patterns, adverse impact, candidate complaints, system drift and recruiter overrides. This lifecycle approach reflects the AI Act's emphasis on risk management and post-market monitoring, but it also adapts those requirements to the organizational context of hiring [4, 19, 32].

The third layer, data and model assurance, governs the technical foundation of hiring decisions. Algorithmic fairness cannot be inferred from the absence of sensitive attributes alone. Research has shown that removing protected characteristics may not prevent discrimination because proxies may remain in the data, prediction targets may be normatively flawed and fairness metrics may trade off against one another [15-16, 27-28]. Data and model assurance therefore requires documented training and validation data, job-relatedness analysis, subgroup performance testing, adverse impact analysis, robustness checks, data minimization and assessment of proxy variables. Documentation mechanisms such as model cards and datasheets can support transparency, but they must be connected to actual decision authority and review routines [17-18].

The fourth layer, human oversight, defines how human judgment is preserved, trained and made consequential. The AI Act requires human oversight for high-risk systems, but oversight may be merely formal if human reviewers lack competence, time, authority or incentives to challenge the

system [4]. In hiring, meaningful oversight requires that a reviewer can understand the system's intended purpose and limitations, access relevant evidence, override outputs, request additional review, document reasons and trigger escalation. Oversight should be located at high-impact decision points such as rejection thresholds, shortlisting cut-offs and final selection recommendations. Human involvement should not be used as a symbolic shield for automated decisions; it should function as a control mechanism that can change outcomes.

The fifth layer, candidate-facing transparency and redress, translates informational and procedural justice into recruitment practice. Candidates should be told when AI materially influences the process, what function the system performs, what information is used, what human review is available and how errors can be challenged [33]. Explanations do not need to reveal proprietary code, but they should be useful enough to support understanding and contestability. Research on explainability emphasizes that explanations are social and contrastive: people want to know why one outcome occurred rather than another, and what could be done differently [34-35]. In hiring, this means that generic statements such as 'an automated tool was used' are insufficient for perceived fairness.

The sixth layer, continuous monitoring and learning, ensures that fairness is not assumed at launch. Hiring systems operate in changing labor markets, changing organizational contexts and changing applicant populations. A model that appears acceptable during validation may create unfair patterns when vacancies, applicant pools or organizational priorities shift. Continuous monitoring should include periodic audits, log review, adverse impact analysis, candidate-experience monitoring, recruiter feedback, incident reporting and revision of thresholds or usage rules. This layer is also where governance becomes adaptive: lessons from complaints, audits and human overrides feed back into procurement, data controls, explanations and oversight.

Table 1 summarizes how key AI Act requirements can be translated into governance functions for algorithmic hiring. The mapping does not treat the AI Act as a complete fairness theory; it shows how regulatory obligations can be connected to justice-relevant organizational controls.

Table 1
 Mapping AI Act requirements to governance functions in algorithmic hiring

AI Act requirement	Recruitment risk addressed	Governance function	Justice contribution
Risk management	Unidentified harms in sourcing, screening, ranking or assessment	Map intended purpose, affected groups, risk scenarios, mitigation measures and residual risks before deployment	Procedural justice through consistency, bias suppression and correctability
Data and data governance	Historical bias, proxy discrimination and non-representative applicant data	Assess data provenance, relevance, representativeness, quality, protected proxies and subgroup performance	Distributive justice through fairer access and reduced adverse impact
Technical documentation	Opaque vendor claims and weak organizational accountability	Require documentation of intended purpose, model logic, validation evidence, limitations and usage conditions	Informational justice and accountability
Record-keeping and logs	Lack of evidence for audits, appeals or incident investigation	Retain logs of system outputs, human overrides, thresholds, errors and complaints where legally permitted	Procedural justice through traceability and correctability
Transparency to deployers and affected persons	Candidates do not know when or how AI influences the process	Provide plain-language notices, role-of-AI explanations and criteria-level information	Informational justice and trust

Table 1 Continued

Human oversight	Automation bias, rubber-stamping and unchallengeable rejection	Assign competent reviewers with authority to intervene, override, escalate and document reasons	Procedural and interpersonal justice
Accuracy, robustness and cybersecurity	Unreliable recommendations, model drift and data-security harms	Test validity, error rates, drift, robustness, cyber risks and failure modes across recruitment contexts	Distributive and procedural justice
Deployer obligations	Provider compliance does not translate into fair organizational use	Define employer-side instructions, monitoring, worker information, log retention and incident procedures	Accountability and legitimacy

Table 2 further specifies the six architectural layers as decision rights and control responsibilities. The table emphasizes that fair algorithmic hiring is a distributed governance task rather than the responsibility of a single HR manager, data scientist or vendor.

Table 2
 Governance layers, decision rights and control responsibilities

Layer	Core decision rights	Key controls	Primary organizational owners
1. Strategic accountability	Which hiring decisions may AI influence and under what fairness commitments?	AI hiring charter, role allocation, vendor requirements, escalation rules	Senior management, HR leadership, legal/compliance
2. Lifecycle risk management	When may a system be procured, deployed, suspended or redesigned?	Pre-use assessment, risk register, mitigation plan, deployment gate, incident response	HR, risk/compliance, data protection officer, procurement
3. Data and model assurance	Which data, targets, features and thresholds are acceptable?	Data documentation, job-relatedness review, subgroup testing, drift analysis, audit support	Data science, HR analytics, vendor, legal
4. Human oversight	Who can challenge, override or escalate algorithmic outputs?	Competence requirements, review protocols, override logs, anti-rubber-stamping safeguards	Recruiters, hiring managers, HR governance board
5. Transparency and redress	What do candidates know and how can they correct errors or seek review?	AI notices, explanations, correction channels, appeal mechanism, respectful communication	HR operations, candidate experience, legal
6. Monitoring and learning	How are outcomes, complaints and system changes converted into improvement?	Periodic audits, adverse impact review, complaint analytics, post-deployment updates	HR governance board, internal audit, compliance

The architecture gives rise to seven theoretical propositions. These propositions articulate theory-based expectations about how governance mechanisms shape fairness and legitimacy in algorithmic hiring and provide a basis for future empirical research and organizational assessment.

Proposition 1. Lifecycle integration proposition. Algorithmic hiring systems are more likely to produce fair and legitimate outcomes when AI Act requirements are embedded into lifecycle governance routines rather than treated as static compliance documents.

Proposition 2. Procedural justice proposition. The relationship between algorithmic hiring and applicant trust is mediated by procedural justice mechanisms, especially consistency, job-relatedness, voice, correctability and the availability of human review.

Proposition 3. Human oversight proposition. Human oversight improves fairness only when reviewers have sufficient competence, authority, time and accountability to challenge system outputs; symbolic or constrained oversight does not produce meaningful justice benefits.

Proposition 4. Data-governance proposition. Representative and well-documented data reduce distributive unfairness only when combined with job-related target definitions, subgroup performance testing and monitoring for proxy discrimination.

Proposition 5. Transparency-redress proposition. Candidate-facing explanations increase perceived informational and procedural justice when they specify the role of AI, the criteria relevant to the decision and the available routes for correction or appeal.

Proposition 6. Role-allocation proposition. Clear allocation of provider, deployer, HR, legal, data protection and managerial responsibilities reduce governance gaps and strengthens accountability for high-risk recruitment decisions.

Proposition 7. Adaptive fairness proposition. Fairness in algorithmic hiring decays over time unless monitoring results, applicant feedback, incidents and human overrides are converted into system updates and governance revisions.

4. Discussion and implications

The proposed architecture makes four theoretical contributions. First, it moves beyond the view that algorithmic hiring fairness is primarily a technical property of the model. Technical controls are necessary, but they are insufficient because hiring fairness also depends on the structure of the decision process, the quality of explanations, the treatment of candidates and the availability of correction. This extends algorithmic fairness research by locating fairness metrics inside organizational decision systems rather than treating them as stand-alone solutions [15-16, 27-28].

Second, the article reframes the EU AI Act as a governance catalyst rather than merely a compliance constraint. The Act's requirements for risk management, data governance, documentation, logging, transparency, human oversight and monitoring can be interpreted as design inputs for fair recruitment governance [4]. However, the Act does not fully specify how these obligations should be translated into HR routines, candidate-facing processes or organizational justice mechanisms. The architecture developed here fills that gap by connecting legal obligations to decision rights and justice outcomes.

Third, the article extends organizational justice theory to high-risk AI-mediated hiring. Traditional selection justice models were developed for human-administered procedures, but their core dimensions remain highly relevant. Algorithmic hiring changes the objects of fairness judgments: candidates may evaluate not only the recruiter or organization but also the perceived legitimacy of the system, the quality of human oversight and the intelligibility of automated recommendations. The architecture therefore adapts distributive, procedural, interpersonal and informational justice to a socio-technical context in which humans and algorithms jointly produce selection outcomes [20-23].

Fourth, the article contributes to HRM governance by specifying how accountability should be distributed across organizational actors, thereby extending broader work on accountability in algorithmic decision-making to the recruitment context [36]. Vendor documentation alone cannot guarantee fair use, because employers configure systems, define job-related criteria, set decision thresholds, determine the weight assigned to algorithmic recommendations, and decide whether candidates may contest outcomes. Conversely, employers cannot govern these systems fairly unless providers supply sufficient documentation, validation evidence, and clear disclosure of system

limitations. Fair algorithmic hiring therefore requires coordination between providers and deployers, internal cross-functional governance, and candidate-facing accountability mechanisms [1-3, 10].

Several practical implications follow. Employers should begin with a recruitment AI inventory that identifies where AI systems materially influence candidate access to work. Each system should be mapped to recruitment stages, affected groups, data inputs, decision outputs, human decision rights and appeal mechanisms. Systems that materially influence screening, ranking or assessment should be governed as high-impact decision systems even when legal classification remains uncertain. This conservative approach is appropriate because reputational, ethical and discrimination risks may arise before formal enforcement action.

Procurement practices should also change. Employers should not purchase algorithmic hiring tools solely on efficiency claims. Vendor contracts should require information about intended purpose, validation evidence, training and testing data, known limitations, subgroup performance, logging functions, audit support, cybersecurity measures and processes for incident response. If a vendor cannot explain how the system was validated for the relevant job context, the employer should not assume that the tool is job-related or fair. This is consistent with the broader accountability literature, which warns that claims of technical neutrality often conceal untested normative choices [2, 15, 19].

HR professionals require new competencies. They do not need to become machine-learning engineers, but they must understand how algorithmic hiring systems structure decisions, what fairness metrics can and cannot show, when human oversight is meaningful, and how to interpret documentation. AI literacy should therefore be role-specific. Recruiters need practical knowledge of system use, overrides and candidate communication. HR leaders need knowledge of procurement, governance and accountability. Legal and compliance teams need knowledge of AI Act obligations, data protection and anti-discrimination law. Managers need knowledge of the limits of algorithmic recommendations.

Candidate-facing governance deserves particular attention. Many organizations focus on internal audits while neglecting the applicant's experience. Yet applicants evaluate fairness through the information they receive and the way they are treated throughout the process. Candidate notices should explain the role of artificial intelligence in clear and accessible language, identify the stage of the recruitment process affected, specify whether human review is available, describe the types of information considered, and provide a mechanism for correcting errors or requesting review. Even when a candidate is not selected, respectful and informative communication can protect organizational legitimacy and reduce perceptions of dehumanization [12, 20, 23].

For policymakers and regulators, architecture suggests that compliance guidance should focus on organizational use cases rather than only technical system categories. Recruitment is a pipeline, and risk arises from the interaction between model outputs, job criteria, recruiter behavior and candidate rights. Regulators can support fair implementation by providing sector-specific templates for fundamental rights impact assessment, examples of meaningful human oversight, documentation standards for recruitment tools and guidance on candidate information notices.

The proposed architecture is subject to several boundary conditions. It is designed for recruitment systems that materially influence access to work and are likely to fall within high-risk or high-impact categories. It is less directly applicable to low-risk administrative tools, such as calendar scheduling systems, provided that these tools do not affect candidate evaluation. Architecture assumes a European Union or European Union-facing organizational context, although many of its elements remain relevant beyond the European Union because fairness, transparency, human oversight, and accountability are widely recognized governance principles across jurisdictions [37-

39]. Nevertheless, the applicable legal duties vary across countries, and organizations should adapt the model to local anti-discrimination, labor, data protection, and collective consultation rules.

Future empirical research should test the propositions in multiple ways. Survey experiments could examine how explanation, appeal and human oversight affect applicants' perceived justice and willingness to continue in a hiring process. Field studies could compare organizations with different governance architectures and measure adverse impact, candidate complaints, time-to-hire and employer attractiveness. Qualitative studies could explore how recruiters interpret algorithmic recommendations and when they feel empowered to override them. Audit studies could examine whether data documentation and subgroup testing reduce discriminatory patterns. These studies would help move the field from conceptual governance design to evidence-based governance evaluation.

5. Conclusions

Algorithmic hiring creates a governance challenge because it reorganizes employment decisions around data, models, vendors, recruiters and automated recommendations. The EU AI Act increases the urgency of this challenge by treating employment-related AI systems as high-risk when they materially influence recruitment, selection and work-related decisions. Yet compliance alone is not sufficient. A hiring system may satisfy formal documentation requirements and still fail to provide candidates with dignity, voice, explanation or meaningful human review.

This article developed a multi-layer decision architecture for fair algorithmic hiring under the EU AI Act. The architecture integrates strategic accountability, lifecycle risk management, data and model assurance, human oversight, candidate-facing transparency and redress, and continuous monitoring and learning. It connects the legal requirements of high-risk AI systems with organizational justice theory and algorithmic fairness research. The central claim is that fair algorithmic hiring is not a property that can be certified once and then assumed. It is an organizational capability that must be designed, enacted, monitored and revised across the recruitment pipeline.

The article's main conclusion is that organizations should shift from tool-centered compliance to decision-centered governance. The relevant question is not only whether an AI tool has been audited, but whether the recruitment decision system is accountable, job-related, explainable, contestable and respectful. When these conditions are met, algorithmic hiring may support more consistent and scalable recruitment. When they are absent, AI may amplify exclusion while appearing neutral. The governance architecture proposed here offers a conceptual foundation for organizations and researchers seeking to build algorithmic hiring systems that are not only efficient but also legitimate and fair.

Acknowledgement

This research was not funded by any grant.

Conflicts of Interest

The authors declare no conflicts of interest.

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this work, the author used ChatGPT as an AI-assisted tool to review the manuscript for consistency of American English usage, grammar, spelling, and minor typographical errors. This tool was not used to generate scientific content, analyses, interpretations, or conclusions.

After using this tool, the author reviewed and edited the manuscript as needed and takes full responsibility for the final version.

References

- [1] Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. *California Management Review*, 61(4), 15–42. <https://doi.org/10.1177/0008125619867910>
- [2] Raghavan, M., Barocas, S., Kleinberg, J., & Levy, K. (2020). Mitigating bias in algorithmic hiring: Evaluating claims and practices. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 469–481. <https://doi.org/10.1145/3351095.3372828>
- [3] Köchling, A., & Wehner, M. C. (2020). Discriminated by an algorithm: A systematic review of discrimination and fairness by algorithmic decision-making in the context of HR recruitment and HR development. *Business Research*, 13, 795–848. <https://doi.org/10.1007/s40685-020-00134-w>
- [4] European Parliament and Council of the European Union. (2024). 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). *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- [5] European Commission. (2026). *Timeline for the implementation of the EU AI Act*. AI Act Service Desk. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/eu-ai-act-implementation-timeline>
- [6] Rigotti, C., & Fosch-Villaronga, E. (2024). Fairness, AI & recruitment. *Computer Law & Security Review*, 53, 105966. <https://doi.org/10.1016/j.clsr.2024.105966>
- [7] Drage, E., & Mackereth, K. (2022). Does AI debias recruitment? Race, gender, and AI’s “eradication of difference”. *Philosophy & Technology*, 35, 89. <https://doi.org/10.1007/s13347-022-00543-1>
- [8] Chen, Z. (2023). Ethics and discrimination in artificial intelligence-enabled recruitment practices. *Humanities and Social Sciences Communications*, 10, 567. <https://doi.org/10.1057/s41599-023-02079-x>
- [9] Budhwar, P., Malik, A., De Silva, M. T., & Thevisuthan, P. (2022). Artificial intelligence – challenges and opportunities for international HRM: A review and research agenda. *The International Journal of Human Resource Management*, 33(6), 1065–1097. <https://doi.org/10.1080/09585192.2022.2035161>
- [10] Malik, A., Budhwar, P., Patel, C., & Srikanth, N. R. (2023). Unlocking the value of artificial intelligence in human resource management through AI capability framework. *Human Resource Management Review*, 33(1), 100899. <https://doi.org/10.1016/j.hrmr.2022.100899>
- [11] Horodyski, P. (2023). Applicants’ perception of artificial intelligence in the recruitment process. *Computers in Human Behavior Reports*, 11, 100303. <https://doi.org/10.1016/j.chbr.2023.100303>
- [12] Chun, J. S., De Cremer, D., Oh, E.-J., & Kim, Y. (2024). What algorithmic evaluation fails to deliver: Respectful treatment and individualized consideration. *Scientific Reports*, 14, 25996. <https://doi.org/10.1038/s41598-024-76320-1>
- [13] Köchling, A., Wehner, M. C., & Ruhle, S. A. (2025). This (AI)n’t fair? Employee reactions to artificial intelligence (AI) in career development systems. *Review of Managerial Science*, 19, 1195–1228. <https://doi.org/10.1007/s11846-024-00789-3>
- [14] Langer, M., König, C. J., & Krause, K. (2017). Examining digital interviews for personnel selection: Applicant reactions and interviewer ratings. *International Journal of Selection and Assessment*, 25(4), 371–382. <https://doi.org/10.1111/ijsa.12191>
- [15] Barocas, S., & Selbst, A. D. (2016). Big data’s disparate impact. *California Law Review*, 104(3), 671–732. <https://doi.org/10.15779/Z38BG31>
- [16] Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59–68. <https://doi.org/10.1145/3287560.3287598>
- [17] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. <https://doi.org/10.1145/3287560.3287596>
- [18] Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- [19] Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44. <https://doi.org/10.1145/3351095.3372873>
- [20] Colquitt, J. A. (2001). On the dimensionality of organizational justice: A construct validation of a measure. *Journal of Applied Psychology*, 86(3), 386–400. <https://doi.org/10.1037/0021-9010.86.3.386>

- [21] Gilliland, S. W. (1993). The perceived fairness of selection systems: An organizational justice perspective. *Academy of Management Review*, 18(4), 694–734. <https://doi.org/10.5465/amr.1993.9402210155>
- [22] Bauer, T. N., Truxillo, D. M., Sanchez, R. J., Craig, J. M., Ferrara, P., & Campion, M. A. (2001). Applicant reactions to selection: Development of the Selection Procedural Justice Scale (SPJS). *Personnel Psychology*, 54(2), 387–419. <https://doi.org/10.1111/j.1744-6570.2001.tb00097.x>
- [23] Hausknecht, J. P., Day, D. V., & Thomas, S. C. (2004). Applicant reactions to selection procedures: An updated model and meta-analysis. *Personnel Psychology*, 57(3), 639–683. <https://doi.org/10.1111/j.1744-6570.2004.00003.x>
- [24] Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10, 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- [25] Whetten, D. A. (1989). What constitutes a theoretical contribution? *Academy of Management Review*, 14(4), 490–495. <https://doi.org/10.2307/258554>
- [26] MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. *Journal of Marketing*, 75(4), 136–154. <https://doi.org/10.1509/jmkg.75.4.136>
- [27] Chouldechova, A. (2017). Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big Data*, 5(2), 153–163. <https://doi.org/10.1089/big.2016.0047>
- [28] Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017). Inherent trade-offs in the fair determination of risk scores. *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*, 67, 43:1–43:23. <https://doi.org/10.4230/LIPIcs.ITCS.2017.43>
- [29] Allal-Chérif, O., Yela Aránega, A., & Castaño Sánchez, R. (2021). Intelligent recruitment: How to identify, select, and retain talents from around the world using artificial intelligence. *Technological Forecasting and Social Change*, 169, 120822. <https://doi.org/10.1016/j.techfore.2021.120822>
- [30] Black, J. S., & van Esch, P. (2020). AI-enabled recruiting: What is it and how should a manager use it? *Business Horizons*, 63(2), 215–226. <https://doi.org/10.1016/j.bushor.2019.12.001>
- [31] van Esch, P., Black, J. S., & Ferolie, J. (2019). Marketing AI recruitment: The next phase in job application and selection. *Computers in Human Behavior*, 90, 215–222. <https://doi.org/10.1016/j.chb.2018.09.009>
- [32] Madaio, M. A., Stark, L., Vaughan, J. W., & Wallach, H. (2020). Co-designing checklists to understand organizational challenges and opportunities around fairness in AI. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14. <https://doi.org/10.1145/3313831.3376445>
- [33] Mourby, M., Ó Cathaoir, K., & Bierre Collin, C. (2021). Transparency of machine-learning in healthcare: The GDPR & European health law. *Computer Law & Security Review*, 43, 105611. <https://doi.org/10.1016/j.clsr.2021.105611>
- [34] Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1–38. <https://doi.org/10.1016/j.artint.2018.07.007>
- [35] Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99. <https://doi.org/10.1093/idpl/ix005>
- [36] Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56–62. <https://doi.org/10.1145/2844110>
- [37] Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28, 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- [38] Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- [39] Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1, 501–507. <https://doi.org/10.1038/s42256-019-0114-4>