CareerIntel –A Conceptual Framework for an AI-Assisted Recruitment and Job-Readiness Platform

Mr. Mayur Suryawanshi
MCA Student
Dr. D.Y Patil Institute of Computer Applications and Management, Pune, India

Mr. Rehan Bagwan
MCA Student
Dr. D.Y Patil Institute of Computer Applications and Management Pune, India

Mr. Maroof Sayyed
MCA Student
Dr. D.Y Patil Institute of Computer Applications and Management Pune, India

Mrs. Bhagyashree Patil
Guide
Dr. D.Y Patil Institute of Computer Applications and Management Pune, India

Abstract— Recruitment technologies increasingly use automated methods to process resumes, compare candidate profiles with job requirements, and support parts of candidate assessment. However, recruitment support and candidate development are often treated as separate activities. A candidate may receive a screening outcome without receiving a clear explanation of the match, an actionable account of missing competencies, or interview practice connected to the target role. This paper proposes CareerIntel, a conceptual framework that connects resume intelligence, job-description analysis, job matching, skill-gap identification, personalized development guidance, and role-aligned interview preparation within a single candidate-centred workflow. The framework also includes a recruiter-support view based on job-related evidence and a governance layer addressing explainability, fairness, privacy, and human oversight. The paper adopts a structured conceptual approach based on a targeted review of relevant research and authoritative guidance. It describes the framework's components, information flow, candidate matching logic, and design considerations. No software implementation, dataset, survey, or empirical evaluation is claimed. Accordingly, the paper does not report measured accuracy, efficiency, or fairness outcomes. CareerIntel is presented as a design proposition whose usefulness and risks require future prototyping, stakeholder evaluation, and empirical validation.

Index Terms—Artificial intelligence, algorithmic hiring, applicant tracking system, explainable AI, job readiness, natural language processing, resume analysis, skill-gap analysis.

I. Introduction

Digital recruitment platforms have changed how job opportunities are advertised and how applications are received. Employers may need to review substantial numbers of applications, while candidates must interpret job descriptions, present their experience clearly, and prepare for selection processes. These activities create information-processing demands on both sides of recruitment. Automated tools can assist with document parsing, search, ranking, and assessment; however, the presence of automation does not by itself ensure that decisions are accurate, understandable, or fair.

Resume screening is one stage of a wider employment process. A resume contains information about education, skills, projects, work experience, and other evidence of capability. A job description expresses a role's responsibilities and expected competencies. Comparing these documents is not simply a matter of counting identical words. Different terms may refer to similar capabilities, while the same term may have different meanings in different contexts. Statistical text-retrieval methods such as term weighting provide a useful basis for comparing documents [1], while contextual language representations can support the interpretation of meaning beyond exact word overlap [2]. Neither approach removes the need for careful design and validation.

A second concern is the limited connection between screening feedback and candidate development. A candidate may be told that a profile is a weak match without being shown which job requirements are evidenced, which are not evidenced, and what practical steps could improve readiness. Interview preparation may also be delivered through tools that are disconnected from the resume and the particular role. This separation can make feedback generic rather than responsive to the candidate's demonstrated projects, skills, and experience.

A third concern relates to the consequences of automated hiring decisions. Algorithmic systems can reproduce or amplify historical patterns in data and can create unequal outcomes even when protected characteristics are not explicitly used as inputs [3],[4]. The literature therefore cautions against treating automated scores as neutral or self-explanatory. Trustworthy AI guidance emphasizes validity, reliability, transparency, explainability, privacy, accountability, and the management of harmful bias [5].

In response to these concerns, this paper proposes CareerIntel as an integrated conceptual framework for candidate development and recruiter decision support. Its central proposition is that resume analysis, job matching, skill-gap feedback, and interview preparation should share a common representation of the candidate profile and the target role. The same evidence used to explain a match should inform the candidate's development plan and the generation of role-specific practice questions. Recruiter-facing functions are included, but automated ranking is not treated as a final hiring decision.

The paper makes four contributions: (1) it defines an integrated workflow connecting candidate readiness with recruitment support; (2) it presents a modular architecture and a transparent, illustrative matching model; (3) it identifies explainability, fairness, privacy, and human oversight as cross-cutting design requirements; and (4) it outlines a future

evaluation agenda without presenting unperformed experiments as completed research.

II. Objectives

The overall objective is to develop a conceptual framework for an AI-assisted platform that links candidate job-readiness support with job-related recruitment assistance.

- To identify functional requirements for resume analysis, job-description interpretation, candidate–job matching, skill-gap analysis, recommendations, and interview preparation.
- To define how information should flow between candidate-facing and recruiter-facing modules.
- To propose an interpretable matching and feedback approach that distinguishes evidenced competencies from missing or uncertain information.
- To identify governance requirements concerning fairness, explainability, privacy, security, and human review.
- To establish a future research agenda for implementation and empirical evaluation.

The paper is guided by the following research questions:

RQ1: How can resume analysis and job-description matching be connected to candidate development activities within one coherent workflow?

RQ2: How can a proposed platform explain job-match outputs and translate identified competency gaps into actionable recommendations?

RQ3: What technical and governance safeguards should be considered when AI is used to support recruitment-related decisions?

III. Literature Review

A. Artificial Intelligence in Recruitment: Artificial Intelligence (AI) is increasingly applied in recruitment activities such as resume screening, candidate selection, job matching, and interview assessment. Mori et al. [6] systematically reviewed the literature on AI in recruiting and selection from an ethical perspective, covering efficiency, fairness, and individual rights. Dadaboyev et al. [7] examined AI applications in employee recruitment and identified opportunities as well as challenges related to algorithmic bias, transparency, and data quality. These studies indicate that AI can support recruitment processes but should operate under appropriate human supervision.

B. Resume Parsing and Job Matching: Resume parsing extracts candidate information such as skills, education, experience, and certifications from submitted documents. Salton and Buckley [1] discussed term-weighting techniques in information retrieval, which underpin keyword-based document comparison using TF-IDF. However, keyword-based methods may fail to recognize similar meanings expressed through different words. Contextual language models such as BERT [2] provide representations that capture semantic relationships within text and can support more meaningful comparison of candidate profiles and job descriptions.

C. Job Recommendation and Skill-Gap Analysis: Job recommender systems help candidates identify opportunities based on their qualifications, skills, and preferences. Ertuğrul and Bitirim [8] reviewed existing job recommender systems and discussed their approaches, applications, and open challenges. However, recommendations alone may not tell a candidate which competencies a particular role requires. Skill-gap analysis addresses this by identifying differences between candidate profile evidence and job requirements, supporting personalized learning and career-development guidance.

D. Interview Preparation, Fairness, and Responsible AI: AI-based technologies can support interview preparation by generating role-specific questions and identifying relevant technical topics from candidate profiles and job descriptions. Generated questions and feedback still require human review for relevance and accuracy. Köchling and Wehner [9] and Fabris et al. [10] highlight concerns about discrimination, fairness, and bias in algorithmic hiring. Concrete cases illustrate these concerns: a recruiting tool was reportedly withdrawn after it was found to disadvantage women [11], and occupation-classification research shows that gender-related bias can persist even when explicit gender indicators are removed [12]. Broader surveys catalogue sources of bias across the machine-learning pipeline [13].

Explainability is often pursued through post-hoc methods such as LIME [14] and SHAP [15], and taxonomies and open challenges of explainable AI are reviewed in [16]. For recruitment, requirement-level evidence mapping may be a more direct form of explanation than post-hoc feature attribution; this is a design proposition, not an established finding.

Regulatory attention is also increasing. The EU Artificial Intelligence Act classifies AI systems used in recruitment and candidate evaluation as high-risk [17], New York City Local Law 144 requires bias audits and candidate notice for automated employment decision tools [18], and India's Digital Personal Data Protection Act, 2023 governs the processing of personal data [19].

E. Literature Summary: The reviewed studies provide a foundation for resume analysis, semantic job matching, job recommendations, skill development, and responsible AI-based recruitment. However, these research areas also highlight the importance of connecting candidate information and job requirements across different stages of the recruitment process. Based on these findings, CareerIntel proposes a conceptual framework integrating resume analysis, job matching, skill-gap identification, development recommendations, and interview preparation within a shared workflow. The framework emphasizes evidence-based explanations and human

oversight. Its practical effectiveness remains subject to future implementation and empirical evaluation.

IV. Research Gap

The reviewed work indicates that algorithmic hiring is a multi-stage socio-technical process and that its risks cannot be reduced to the performance of a matching algorithm alone [3],[10]. Existing research and tools address individual stages, such as matching, recommendation, or interview practice, but these functions are frequently treated as separate utilities. This paper identifies a practical design opportunity: connect resume screening with candidate-oriented skill development and interview preparation, and define that connection as a conceptual contribution.

Observed design concern	**Implication for a proposed integrated framework**
Resume analysis and job matching may be treated as separate utilities.	Use one normalized candidate profile and one structured representation of role requirements.
A match score may not reveal which evidence influenced it.	Provide requirement-level explanations and identify missing, partial, and uncertain evidence.
Candidate feedback may remain generic.	Map each confirmed skill gap to a role-relevant development action.
Interview practice may be disconnected from the candidate's profile.	Generate practice topics from the target role and candidate's documented evidence.
Recruiter ranking can be mistaken for an objective hiring decision.	Keep outputs decision-support only, with human review and auditability.
Fairness and privacy may be addressed after system design.	Include governance, data minimization, access control, and monitoring throughout the lifecycle.

The research problem is therefore defined as follows: how can an AI-assisted system be conceptually organized so that candidate profile analysis, job-specific matching, actionable readiness guidance, and interview practice operate as connected functions, while recruitment-related outputs remain explainable, privacy-conscious, and subject to human oversight?

The proposed contribution is architectural and integrative. This paper does not claim that no existing platform combines any of these functions. Instead, it proposes a coherent information flow in which the output of one module becomes the grounded input to the next, and in which explanations and governance are designed as shared services rather than optional additions.

V. Research Methodology

A. Research Design: This study adopts a conceptual research design supported by a targeted, structured narrative review of literature and authoritative technical guidance. It is descriptive and analytical rather than experimental. The purpose is to identify relevant concepts, compare design considerations, define a proposed framework, and derive propositions for future validation.

B. Literature Identification Approach: A targeted search was conducted in [list only the databases you really used, e.g., Google Scholar, IEEE Xplore, ACM Digital Library, Scopus] during [month year to month year]. Search terms included [keep only the terms you really used, e.g., “AI recruitment”, “algorithmic hiring fairness”, “resume screening”, “job matching”, “TF-IDF”, “contextual language models”, “explainable AI”, “AI risk management”]. Records were screened by title and abstract, and relevant items were followed to publisher, DOI, or standards-body pages. The search aimed to identify foundational and directly relevant sources; it was not an exhaustive systematic review.

C. Inclusion and Exclusion Principles: Priority was given to peer-reviewed research, established conference proceedings, foundational technical publications, and authoritative institutional guidance that directly informed a component of the proposed framework. Sources were considered relevant when they addressed recruitment automation, text matching, algorithmic fairness, or AI governance. Sources with incomplete bibliographic details or unverifiable claims were excluded. No counts of records identified, screened, or included are reported because a formal screening log was not maintained.

D. Synthesis Procedure: The selected material was organized into five design themes: (1) candidate information extraction, (2) candidate–job matching, (3) skill-gap and development support, (4) interview preparation, and (5) responsible recruitment governance. The themes were compared to identify dependencies between modules and potential failure points. These dependencies informed the proposed architecture, the illustrative matching logic, and the governance requirements described in Sections VI and VII.

E. Scope and Methodological Boundary: No participants were recruited, no questionnaire was administered, no resume dataset was collected, and no prototype was implemented or evaluated for this paper. The study therefore has no sample size, experimental group, performance metrics, or statistical results. The proposed framework and equations are design specifications for future work, not validated algorithms.

VI. Proposed CareerIntel Conceptual Framework

Framework Overview: CareerIntel is conceived as a two-sided platform with a shared intelligence layer. The candidate-facing side helps a user understand how a resume relates to a selected job, identify evidence gaps, plan development activities, and practise role-relevant interview questions. The

recruiter-facing side supports structured review of candidate evidence against published job requirements. Both sides use the same normalized data model and explanation services, while access permissions restrict which candidate information each role may view.

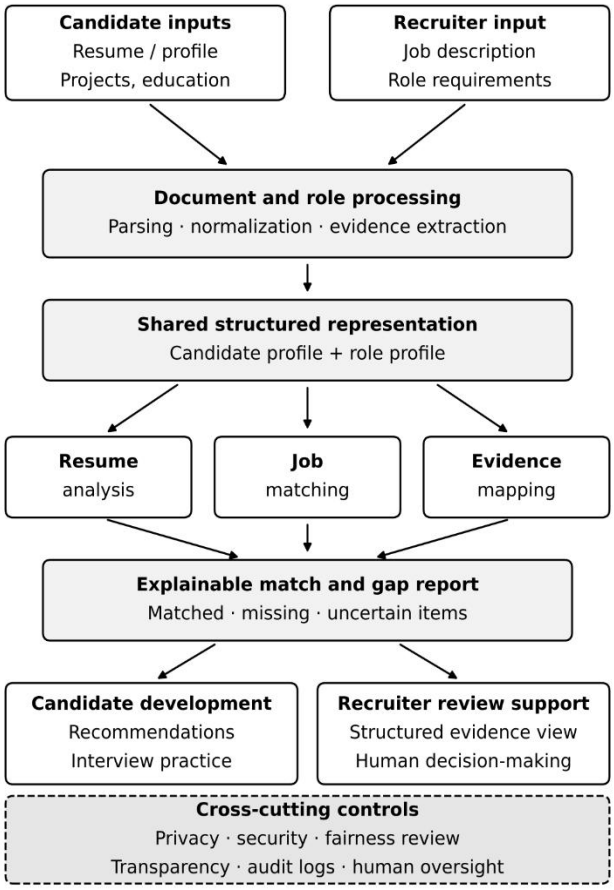


Fig. 1. The diagram represents the proposed information flow; it is not a deployment or implementation diagram.

VII. Technical Design Considerations

A. Data Model: A conceptual data model may contain the following entities: Candidate, ResumeDocument, CandidateEvidence, Skill, EducationRecord, Project, ExperienceRecord, JobRole, JobRequirement, MatchAssessment, SkillGap, Recommendation, InterviewQuestion, RecruiterReview, and AuditEvent. CandidateEvidence should preserve provenance by linking each extracted item to a document segment or candidate confirmation. JobRequirement should store the original wording and any recruiter-assigned priority.

ID	Requirement
FR1	Accept resumes and job descriptions and extract text with document provenance
FR2	Normalize skills, education, dates, and projects, with candidate confirmation of uncertain fields
FR3	Represent each job requirement with its original wording and priority
FR4	Classify each requirement as evidenced, partial, not evidenced, or uncertain
FR5	Generate requirement-level explanations referencing candidate evidence
FR6	Map confirmed gaps to development actions and role-aligned practice questions
FR7	Provide a recruiter view for human review, with override and audit logging

B. Processing Workflow

Upload and validate the resume; extract text while retaining document provenance.

Normalize sections, dates, skill names, and project descriptions; request candidate confirmation for uncertain fields.

Parse the job description into explicit requirements and preserve the original wording.

Compare candidate evidence with each role requirement using a disclosed lexical, semantic, or hybrid method.

Generate requirement-level explanations and identify missing, partial, and uncertain evidence.

Map confirmed gaps to recommendations and create interview-practice topics from the same role profile.

Present recruiter-facing evidence with human review, access controls, and audit records.

C. Illustrative Matching Logic: A future implementation could use a staged matching process. Let $R = \{r_1, \ldots, r_n\}$ be the requirements of a role, each with a priority weight w_i ($\Sigma w_i = 1$). Lexical similarity can use TF-IDF term weights [1]:

$$\text{tfidf}(t,d) = \text{tf}(t,d) \cdot \log(N / \text{df}(t)) \quad (1)$$

and cosine similarity between vectors $\mathbf{v}$ of the requirement and a candidate evidence segment:

$$\text{sim}(q, e) = (v_q \cdot v_e) / (\|v_q\| \|v_e\|) \quad (2)$$

A contextual encoder [2], [20] can provide a semantic similarity S_i between embeddings of the requirement and the evidence. A combined requirement-level score is then:

$$m_i = \alpha \cdot L_i + \beta \cdot S_i + \gamma \cdot T_i, \text{ with } \alpha + \beta + \gamma = 1 \quad (3)$$

where L_i is lexical similarity, S_i is semantic similarity, and T_i is a structured check (for example qualification or experience level) for the best-matching evidence segment. The status of each requirement could be:

$$\text{status}(r_i) = \text{Evidenced if } m_i \geq \tau_h; \text{ Partial if } \tau_l \leq m_i < \tau_h; \text{ Not evidenced otherwise} \quad (4)$$

with "Uncertain" assigned when the evidence comes from an unconfirmed parse. An optional summary indicator could be $\mathbf{Q = \Sigma\, w_i \cdot \kappa_i}$, with $\kappa_i = 1, 0.5$, or 0 for evidenced, partial, or not evidenced (5). The weights α, β, γ and thresholds τ_l, τ_h are open design parameters that would require calibration against

expert-annotated data; the summary indicator is illustrative and should not be used alone for ranking. Component-level outputs and explanations are retained rather than collapsed into an opaque score.

Algorithm 1 :

Input: candidate profile C, job profile J

1. validate_and_normalize(C, J)

2. for each requirement r in J.requirements:

evidence = retrieve_candidate_evidence(C, r)

status = classify_evidence(evidence, r)

explanation = explain(status, evidence, r)

store(r, status, explanation)

3.Gaps=requirements_needing_confirmation_or_developme nt()

4. recommendations = map_gaps_to_learning_actions(gaps)

5. practice_set = generate_role_aligned_questions(C, J)

6. present_candidate_report(recommendations, practice_set, explanations)

7.present_recruiter_view_with_human_review(explanations)

Output: explainable decision-support report

D. Non-Functional Requirements

Requirement	**Conceptual design expectation**
Explainability	Show evidence, requirement mapping, uncertainty, and limitations.
Fairness	Assess data and outcomes for harmful disparities; do not assume bias-free operation.
Privacy and security	Minimize data, restrict access, protect documents, and define retention.
Reliability	Handle parsing failures and missing fields without silently inventing information.
Usability and accessibility	Provide understandable feedback and accessible candidate-facing interactions.
Auditability	Record model/rule versions, output timestamps, access, and human overrides.
Scalability	Separate document processing, matching, and recommendation services where deployment needs justify it.

VIII. Framework Analysis and Discussion

The proposed framework is intended to address a workflow-level problem rather than introduce a claim of algorithmic novelty. Its central design feature is the reuse of a shared candidate profile and role representation across matching, gap analysis, recommendations, and interview practice. This creates traceability: a candidate can see which role requirement led to a gap, which evidence was considered, and why a suggested learning activity or practice question was generated.

A. Candidate-Centred Value Proposition: For students, fresh graduates, and early-career applicants, the framework gives particular attention to academic projects, coursework, certifications, and practical work that may not appear as conventional employment history. This does not mean that academic projects are equivalent to professional experience. Rather, they can be represented as evidence of selected competencies and reviewed in context. Candidate-facing explanations may help users improve the clarity of their resumes and identify areas for further learning.

B. Recruiter-Centred Value Proposition: For recruiters, the proposed system offers structured access to job-related evidence and may help organize large application pools. Its use should remain assistive. Recruiters should not treat a rank, similarity value, or generated summary as a final judgment. Human review is necessary to consider context, correct extraction errors, and assess evidence that the system cannot reliably represent.

C. Why a Single Score Is Insufficient: A composite match score can simplify comparison but may conceal important differences. Two candidates with the same score may have different strengths, evidence quality, or missing information. A high lexical similarity may reflect repeated job-description terms rather than demonstrated competence. A low similarity may arise because the candidate uses different terminology. CareerIntel therefore proposes requirement-level explanations alongside any optional summary indicator.

D. Fairness and Accountability Trade-Offs: A system that aims to improve consistency may still reproduce historical patterns in its data or penalize non-traditional career paths. Removing explicit demographic variables does not eliminate proxy effects. Fairness evaluation should consider the complete recruitment process, including the data used, the role requirements, the threshold or ranking policy, the way recruiters interpret the output, and the impact on candidates. The multidisciplinary literature emphasizes that fairness is context-dependent and that no single metric captures every relevant concern [10].

E. Conceptual Propositions for Future Testing:

P1: A shared candidate and job representation can improve traceability between matching explanations, skill-gap reports, and interview-practice content.

P2: Requirement-level explanations can provide more actionable feedback than an unexplained aggregate match score.

P3: Candidate confirmation of extracted profile information can reduce the risk that parsing errors are treated as genuine competency gaps.

P4: Human review, audit logging, and fairness monitoring are necessary safeguards for recruiter-facing use; their effectiveness must be assessed in the deployment context.

These propositions are conceptual expectations, not findings. They require empirical investigation before any claim can be made about effectiveness.

F. Answering the Research Questions: *RQ1:* Resume analysis and job matching can be connected to development activities through a shared normalized candidate profile and role profile, so that the same evidence supports matching, gap analysis, recommendations, and interview practice. *RQ2:* Match outputs can be explained through requirement-level statuses (evidenced, partial, not evidenced, uncertain) that reference candidate evidence, and each confirmed gap can be mapped to a role-relevant development action. *RQ3:* Safeguards to consider include candidate confirmation of extracted data, human review of recruiter-facing outputs, audit logging, data minimization and access control, fairness monitoring, and compliance review under applicable regulation [17]–[19]. These answers are conceptual and require empirical validation.

IX. Limitations

The paper presents a conceptual framework and does not report a working software prototype or empirical evaluation.

The literature identification process was targeted and narrative; it was not a registered or exhaustive systematic review, and no PRISMA-style screening counts are claimed.

The matching function is specified at a conceptual level. Feature definitions, model selection, weights, thresholds, and calibration remain open design decisions.

Resume parsing may fail on complex layouts, scanned documents, multilingual content, or incomplete information.

Skill names and proficiency levels are difficult to infer reliably from short resume statements; candidate confirmation is therefore important.

Generative interview questions and recommendations may be irrelevant, inaccurate, repetitive, or overly confident.

Fairness cannot be guaranteed by excluding protected attributes alone; proxy effects, data quality, and human use may still create unequal outcomes.

Legal requirements for employment-related AI and personal-data processing vary across jurisdictions and must be assessed before deployment.

X. Future Directions

Future work should proceed in stages. First, researchers should develop a prototype with transparent parsing and matching components. Second, the system should be evaluated using a carefully governed dataset containing diverse resume formats, job descriptions, occupations, and candidate backgrounds. Dataset collection should follow consent, minimization, access, and retention requirements. Third, matching quality should be evaluated against expert-annotated job requirements using appropriate precision, recall, ranking, and calibration measures. Parsing quality should be assessed separately from matching quality so that extraction errors are not confused with model errors.

Fourth, candidate studies should assess whether explanations and recommendations are understandable, relevant, and actionable. Fifth, recruiter studies should examine whether the tool improves review consistency without encouraging automation bias. Sixth, fairness audits should assess subgroup outcomes and error patterns where lawful and ethically justified, while documenting limitations caused by sample size and sensitive-data restrictions. Finally, longitudinal studies may investigate whether candidate use is associated with changes in resume quality, skill development, interview confidence, or employment outcomes. These are proposed research activities, not completed work.

XI. Conclusion

This paper has presented CareerIntel as a conceptual framework for connecting AI-assisted recruitment support with candidate job-readiness development. The framework combines resume parsing, job-description analysis, explainable candidate–job matching, ATS-oriented feedback, skill-gap analysis, personalized recommendations, interview practice, and recruiter decision support. Its principal design proposition is that these functions should share a common candidate and role representation so that feedback remains connected to the evidence used in matching.

The paper also argues that explainability, fairness, privacy, security, and human oversight must be treated as cross-cutting requirements rather than optional additions. A match score should not be interpreted as a definitive measure of ability, and a recruiter-facing ranking should not replace human judgment. Because CareerIntel has not been implemented or empirically evaluated in this study, its benefits remain propositions requiring future validation. The framework provides a basis for subsequent prototype development, stakeholder evaluation, and responsible assessment.

References

[1] G. Salton and C. Buckley, “Term-weighting approaches in automatic text retrieval,” *Inf. Process. Manage.*, vol. 24, no. 5, pp. 513–523, 1988, doi: 10.1016/0306-4573(88)90021-0.

[2] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. NAACL-HLT*, Minneapolis, MN, USA, 2019, pp. 4171–4186, doi: 10.18653/v1/N19-1423.

[3] M. Raghavan, S. Barocas, J. Kleinberg, and K. Levy, “Mitigating bias in algorithmic hiring: Evaluating claims and practices,” in *Proc. FAT* ‘20*, Barcelona, Spain, 2020, pp. 469–481, doi: 10.1145/3351095.3372828.

[4] S. Barocas and A. D. Selbst, “Big data’s disparate impact,” *California Law Rev.*, vol. 104, no. 3, pp. 671–732, 2016, doi: 10.15779/Z38BG31.

[5] E. Tabassi, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, 2023, doi: 10.6028/NIST.AI.100-1.

[6] M. Mori et al., "A systematic literature review on artificial intelligence in recruiting and selection: a matter of ethics," *Personnel Review*, vol. 54, no. 3, pp. 854–878, 2025, doi: 10.1108/PR-03-2023-0257.

[7] S. M. U. Dadaboyev, J. Abdullayeva, N. Abbosova, A. Suleymenova, and K. Mamadjanova, "Role of artificial intelligence in employee recruitment: systematic review and future research directions," *Discover Global Society*, vol. 3, Art. no. 99, 2025, doi: 10.1007/s44282-025-00246-w.

[8] D. Ç. Ertuğrul and S. Bitirim, "Job recommender systems: a systematic literature review, applications, open issues, and challenges," *J. Big Data*, vol. 12, Art. no. 140, 2025, doi: 10.1186/s40537-025-01173-y.

[9] A. Köchling and M. C. Wehner, "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*, vol. 13, pp. 795–848, 2020, doi: 10.1007/s40685-020-00134-w.

[10] A. Fabris et al., "Fairness and bias in algorithmic hiring: A multidisciplinary survey," *ACM Trans. Intell. Syst. Technol.*, vol. 16, no. 1, Art. no. 16, 2025, doi: 10.1145/3696457.

[11] J. Dastin, "Insight—Amazon scraps secret AI recruiting tool that showed bias against women," Reuters, Oct. 10, 2018. [Online]. Available: [add exact URL]. Accessed: [date].

[12] M. De-Arteaga et al., "Bias in bios: A case study of semantic representation bias in a high-stakes setting," in *Proc. FAT* '19*, Atlanta, GA, USA, 2019, pp. 120–128, doi: 10.1145/3287560.3287572.

[13] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan, "A survey on bias and fairness in machine learning," *ACM Comput. Surv.*, vol. 54, no. 6, Art. no. 115, 2021, doi: 10.1145/3457607.

[14] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, San Francisco, CA, USA, 2016, pp. 1135–1144, doi: 10.1145/2939672.2939778.

[15] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 4765–4774.

[16] A. Barredo Arrieta et al., "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Inf. Fusion*, vol. 58, pp. 82–115, 2020, doi: 10.1016/j.inffus.2019.12.012.

[17] Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), *Official J. Eur. Union*, 2024.

[18] New York City Council, "Local Law 144 of 2021: Automated employment decision tools," N.Y.C. Admin. Code § 20-870 et seq.

[19] Government of India, *The Digital Personal Data Protection Act, 2023*, Act No. 22 of 2023.

[20] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in *Proc. EMNLP-IJCNLP*, Hong Kong, China, 2019, pp. 3982–3992, doi: 10.18653/v1/D19-1410.