

Intelligent Screening of Online Job Postings for Recruitment Risk and Responsible Hiring Support

Smart Online Job Postings for Recruitment Risk

Tilottama Azad

*#Associate Professor Department of management, TIT College, India,
Bhopal*

drtilottamaa@gmail.com

Abstract— Artificial intelligence (AI) is increasingly transforming recruitment and talent-acquisition processes; however, concerns regarding opaque predictions, unreliable confidence estimates, algorithmic bias, and inadequate human oversight continue to limit its responsible adoption. This study proposes an explainable and confidence-adaptive AI framework for identifying potentially problematic recruitment patterns in online job postings while retaining human control over uncertain decisions. A publicly available corpus of 22,000 U.S. job advertisements was analyzed, of which 21,879 records remained after quality screening. Weak-supervision rules were developed around four recruitment-risk dimensions: role alignment, experience rationality, workload feasibility, and internal discrepancy. Five machine-learning models—Logistic Regression, Support Vector Machine, Random Forest, XGBoost, and LightGBM—were compared using a leakage-controlled evaluation protocol. LightGBM achieved the strongest validation performance and was subsequently calibrated using sigmoid calibration with a validation-optimized decision threshold. On the held-out test set, the final model achieved 87.08% accuracy, 87.25% precision, 97.16% recall, 91.94% F1-score, 0.621 Matthews correlation coefficient, 0.939 ROC-AUC, and 0.980 PR-AUC. A confidence-adaptive Human-in-the-Loop mechanism routed 44.98% of cases for review while capturing 96.30% of observed model errors. SHAP explanations, multidimensional risk profiling, semantic-hybrid representation, and potential bias-trigger screening further improved interpretability and auditability. The proposed system is positioned as recruitment-content decision support rather than an autonomous hiring mechanism.

Keywords— Artificial intelligence, recruitment analytics, human resource management, explainable AI, human-in-the-loop, LightGBM, responsible AI, job-posting analysis

I. INTRODUCTION

Artificial intelligence (AI) has evolved from a novelty into a critical tool in human resource management (HRM). AI-based recruitment systems can analyze vast amounts of data, identify patterns, and enhance the efficiency and consistency of screening processes. However, recruitment is a sensitive area where AI recommendations can impact individuals' employment opportunities and the overall workforce. Thus, besides being accurate, recruitment AI must be capable of providing evidence, confidence levels, human supervision, and a set of limitations to explain its results.

The importance of AI-enabled HRM was emphasized by previous research. Studies have reported strong relationships

between recruitment AI and improved HR practices, decision-making, and sustainable HR performance [1]. Explainable HRM AI was linked to higher levels of applicants' perception of transparency, organizational support, fairness, and innovativeness [2]. Organizational resources, human capability, technology, and dynamic capabilities were found to be vital areas for successfully implementing AI-based recruitment [3]. At the individual level, the focus of prior research was on the relationships between recruiting AI and employees' perceived usefulness, trust, and threat as well as their AI aversion [4]. The impact of AI applications was also found to be positive in terms of improving the efficiency and agility of HR decision-making [5]. Despite such benefits, wider people-analytics research highlights concerns about algorithmic bias, privacy, employee trust, and excessive automation [6]. Meanwhile, implementation issues related to integration, skills, and the need for human-centered decision making emerge around AI-driven HR solutions [7]. Explainability is particularly important in recruitment because AI recommendations could influence human hiring decisions, while the effects of XAI may vary across candidate contexts [8]. Ethical HRM research therefore focuses on intelligibility, accountability, and responsible algorithmic decision making [9]. Human-resource resilience studies further show that AI should enhance the capability of an organization rather than replace professional expertise [10]. Human-AI collaboration studies similarly emphasize the importance of technological capability and human judgments [11]. Sustainability-oriented HRM research also reveals that governance, HR-IT collaboration, and workforce AI literacy influence the performance of AI-enabled systems [12]. Research on HR-system design also underscores the importance of explainability, task complexity, social acceptability, and professional identity in implementing AI [13]. Cross-functional AI studies also identify governance, organizational integration, and workforce readiness as enablers for AI-related performance gains [14]. Reviews on responsible AI in HRM conclude that fairness, accountability, transparency, and meaningful human control are vital for trustworthy AI deployment [15].

Collectively, such research highlights the importance of AI adoption, explainability, ethics, and human involvement, which are largely examined separately. There is therefore a need for an operational recruitment-support system that determines what evidence drives a prediction, how reliable it is, which recruitment-risk dimension it involves, and when a human

review is required. In this regard, the present study develops an explainable and confidence-adaptive AI framework that integrates weak supervision, machine-learning benchmarking, probability calibration, SHAP explanations, multidimensional risk profiling, semantic representation, bias-trigger screening, and Human-in-the-Loop review.

II. METHODOLOGY

2.1 Dataset and Analytical Design

The study made use of a publicly available dataset of 22,000 United States-based job advertisements. Initial screening of the data quality resulted in the removal of 121 observations with technical or insufficient content, leaving 21,879 eligible job postings. A total of 600 postings were included in the study's evaluation activities, with 150 of them used in a double-annotation experiment. In the main supervised experiment, 2,089 records were included, of which 1,587 were labeled positively and 502 negatively by weak supervision, while the final held-out test set accounted for 418 cases.




Fig. 1. Overall Architecture of the Proposed Explainable and Confidence-Adaptive AI Recruitment-Risk Framework

Figure 1 represents the whole workflow, from job-posting ingestion till final interpretable recruitment-support evidence. Both structured and textual information are transformed into predictive features and assessed within multiple machine-learning models. The chosen model is calibrated ahead of the SHAP explanation, confidence estimation, multidimensional profiling, and human-review routing. Potential linguistic warning signals are still considered separate review layer rather than proof of discriminatory behavior.

2.2 Weak-Supervision Risk Construction

Because independently verified labels were unavailable for the complete dataset, interpretable weak-supervision rules were developed around four dimensions: role alignment, experience rationality, workload feasibility, and internal discrepancy.

For dimension d , labeling function l applied to posting i was expressed as

$$\lambda_{ij}^{(d)} \in \{-1, 0, 1\}, \quad (1)$$

where -1 denotes abstention and 0/1 indicate alternative provisional outcomes. Rules were permitted to abstain instead of assigning a class where textual evidence was insufficient.

The four-dimensional recruitment-risk representation was defined as

$$R_i = [R_i^{role}, R_i^{exp}, R_i^{work}, R_i^{disc}]. \quad (2)$$

The supervised predictive analysis principally focused on the resolved workload-related risk signal. Model performance must therefore be interpreted as predictive agreement with weakly supervised labels rather than independently verified ground truth.

2.3 Feature Engineering and Machine-Learning Models

A set of 14 interpretable handcrafted features was used in the primary experiment, where leakage was controlled. This set includes document structure, title-related features, requirement intensity, experience references, seniority indicators, educational information, numerical patterns, sensitive-language mentions, and formatting-related features.




Fig. 2. Dataset Reduction and Analytical Flow from the Original 22,000 Postings to the 21,879-Record Research Corpus and Held-Out Evaluation Data

The feature vector was represented as

$$x_i = [x_{i1}, x_{i2}, \dots, x_{i14}]^T. \quad (3)$$

The learning objective was to estimate

$$\hat{p}_i = P(y_i = 1 | x_i). \quad (4)$$

Five models were evaluated: Logistic Regression (LR), Support Vector Machine (SVM), Random Forest (RF), XGBoost, and LightGBM. The models were assessed using accuracy, precision, recall, F1-score, MCC, ROC-AUC, and PR-AUC rather than accuracy alone.

Figure 2 distinguishes corpus-level analysis from supervised predictive evaluation. This separation is important because the complete corpus supports descriptive risk profiling and linguistic screening, whereas only observations with resolved provisional labels support supervised learning. The held-out test data remained separate from model and threshold selection.

2.4 Calibration, Thresholding and Explainability

LightGBM was selected from the comparative validation experiment. Because raw model probabilities may not represent reliable confidence levels, sigmoid calibration was subsequently applied:

$$\tilde{p}_i = \frac{1}{1 + \exp(Af(x_i) + B)}. \quad (5)$$

The decision threshold was determined from validation data using MCC as the optimization criterion:

$$\tau^* = 0.44. \quad (6)$$

The final classification decision was

$$\hat{y}_i = \begin{cases} 1, & \tilde{p}_i \geq 0.44, \\ 0, & \tilde{p}_i < 0.44. \end{cases} \quad (7)$$

SHAP was used for model interpretation:

$$f(x_i) = \phi_0 + \sum_{j=1}^p \phi_{ij}, \quad (8)$$

where ϕ_{ij} quantifies the contribution of feature j to a particular prediction.

2.5 Confidence-Adaptive Human Oversight

A standardized certainty score was derived from the calibrated probability:

$$Q_i = |2\tilde{p}_i - 1|. \quad (9)$$

Lower values correspond to predictions closer to the decision boundary. Cases were subsequently allocated to mandatory human review, standard review, or low-intervention categories.

2.6 Semantic-Hybrid Extension and Potential Bias Screening

A secondary representation experiment incorporated semantic information from complete job descriptions. TF-IDF produced up to 30,000 textual dimensions, which were reduced using truncated SVD to 100 semantic components. The

semantic vector s_i was combined with the 14 handcrafted variables:

$$h_i = s_i \oplus x_i \in \mathbb{R}^{114}. \quad (10)$$

In addition, a narrow linguistic screening layer detected potential age-preference and exclusionary personality wording. These indicators were treated as review triggers rather than evidence of demographic discrimination.

	○ Not primary capability	● Implemented capability					
Baseline ML	●	●	●	○	○	○	○
Probability calibration	●	●	○	○	○	○	○
Threshold optimization	●	●	●	○	○	○	○
SHAP explainability	●	●	○	●	○	○	○
Bias-trigger screening	○	○	○	●	●	●	●
Adaptive HITL	●	●	●	●	●	●	●
Evidence & recommendation	○	●	●	●	●	●	●
	Prediction	Probability	Decision rule	Explainability	Bias screening	Human review	Recommendation

Fig. 3. Capability Matrix of the Proposed AI-Assisted Recruitment-Risk Assessment Framework

Figure 3 summarizes how the principal components of the proposed framework contribute to prediction, probability estimation, decision rules, explainability, bias screening, human review, and recommendation generation. Baseline machine-learning components primarily support prediction, whereas calibration and threshold optimization strengthen probability reliability and decision control. SHAP and bias-trigger modules extend the framework toward interpretation and responsible screening, while adaptive HITL integrates all major capabilities. The matrix therefore highlights the progressive transition from conventional prediction toward explainable and human-centered recruitment decision support

III. RESULTS AND DISCUSSION

3.1 Model Benchmarking

The five machine-learning models demonstrated substantial differences in predictive performance. Table 1 indicates a progressive improvement from linear and kernel-based classifiers toward tree-based ensemble models. LightGBM achieved the highest validation accuracy, F1-score, MCC, ROC-AUC, and PR-AUC, while Random Forest remained closely competitive. The improvement in MCC from 0.561 for Logistic Regression to 0.725 for LightGBM is particularly important because MCC incorporates all four cells of the confusion matrix and is more informative than accuracy under class imbalance.

TABLE 1. VALIDATION PERFORMANCE OF THE FIVE CANDIDATE MODELS

Model	Accu racy	Preci sion	Rec all	F1- score	M CC	R OC- AUC	P R- AUC

Logistic Regression	84.69 %	88.25 %	92.1 4%	90.1 5%	0. 561	0. 926	0. 976
SVM	87.56 %	89.58 %	94.6 5%	92.0 5%	0. 641	0. 915	0. 964
XGBoost	89.47 %	92.81 %	93.4 0%	93.1 0%	0. 709	0. 947	0. 982
Random Forest	89.95 %	92.86 %	94.0 3%	93.4 4%	0. 720	0. 948	0. 982
LightGBM	90.19 %	92.62 %	94.6 5%	93.6 2%	0. 725	0. 950	0. 983




Fig. 4. Comparative Performance of Logistic Regression, SVM, XGBoost, Random Forest, and LightGBM Across Validation Metrics

Figure 4 demonstrates that LightGBM selection was supported by multiple evaluation criteria rather than a single accuracy score. The clustering of XGBoost, Random Forest, and LightGBM at the upper end of the benchmark suggests that nonlinear interactions among recruitment-text characteristics are important. LightGBM was therefore retained for calibration and final held-out evaluation.

3.2 Final Test Performance

Table 2 shows that calibration and threshold optimization perform different functions. Calibration improves probability reliability, whereas threshold adjustment changes the sensitivity–specificity trade-off. The 0.44 threshold prioritizes

recall and reduces false negatives rather than maximizing every metric simultaneously.

The final LightGBM model achieved an accuracy of 87.08%, precision of 87.25%, recall of 97.16%, F1 -score of 91.94%, MCC of 62.1%, ROC-AUC of 93.9%, and PR-AUC of 98%. Only 9 patients were falsely classified as negative by the model. Thus, the screening method demonstrated high performance, but the number of false positives is significant enough to warrant further calibration of thresholds and human assessment instead of relying solely on automatic decisions.

3.3 Calibration and Threshold Analysis

Probability calibration improved reliability while preserving discrimination. The validation Brier score decreased from approximately 0.0740 to 0.0722 after sigmoid calibration. The final test Brier score was 0.0857.

Threshold optimization produced an operating point of 0.44, selected using validation MCC. This threshold was intentionally chosen before final test evaluation and increased sensitivity toward positive cases.

TABLE 2. EFFECT OF CALIBRATION AND DECISION-THRESHOLD OPTIMIZATION

Model stage	Threshold	F1	Recall	MCC	Brier Score	Interpretation
Raw LightGBM	0.50	0.921	Lower than final	0.653	0.0943	Uncalibrated baseline
Calibrated LightGBM	0.50	Improved probability reliability	—	—	0.0857	Better probability calibration
Final operating model	0.44	0.919	0.972	0.621	0.0857	Recall-oriented screening point




Fig. 5. Probability Calibration and Validation-Based Threshold Trade-Off of the Selected LightGBM Classifier

Figure 5 demonstrates the need for distinguishing between probability quality and the screening threshold. The former aids in understanding the prediction's probability, whereas the latter makes sense of that probability through the use of a screening decision. This is particularly critical as confidence decides the effort of human review.

3.4 Explainability

According to the global SHAP analysis, the sentence count accounted for about 26.16% of the mean absolute SHAP importance, followed by the requirement indicators with 22.63%, and sensitive language mentions with 10.39%. Some other important predictors are the attributes of titles, average sentence length, number density, seniority indicators, capitalization, and education attributes. Figure 6 depicts how each of the individual job posting attributes affect the recruitment risk prediction probability in the 418 observation test set. The variables of sentence count and requirement indicators have the widest SHAP distribution and thus are the most impactful on model predictions. The sensitive language mentions and attributes of titles also play a significant role. Positive SHAP values increase the predictions, and negative values lower the prediction probabilities.




Fig. 6. SHAP Impact Distribution of Recruitment-Risk Predictions Across the Held-Out Test Set.

3.5 Confidence-Adaptive Human-in-the-Loop Performance

Among 418 test cases, 99 required mandatory review, 89 standard review, and 230 low intervention. Thus, 188 cases (44.98%) received substantive human review. Of 54 total errors, 52 occurred within the review groups and only 2 outside them, giving an error-capture rate of:

$$ECR = \frac{52}{54} \times 100 = 96.30\%$$

This confirms that confidence-adaptive routing concentrated nearly all observed errors within less than half of the evaluated cases.

Figure 7 compares the allocation of 418 test cases across mandatory review, standard review, and low-intervention categories with the distribution of 54 observed prediction errors. Mandatory and standard review together account for 188 cases (44.98%), yet capture 52 of 54 errors (96.30%). Specifically, 72.2% of errors occur in mandatory-review cases, 24.1% in standard-review cases, and only 3.7% in low-intervention cases. This concentration supports selective human oversight by directing review effort toward cases where model uncertainty and observed error risk are greatest.



Fig. 7. Confidence-Adaptive Human-in-the-Loop Case Allocation and Retrospective Concentration of Prediction Errors Across Review Levels.

3.6 Multidimensional Recruitment-Risk Profile

The full 21,879-posting corpus was subsequently characterized across four recruitment-risk dimensions. Overall scores classified 4,115 postings as low risk, 6,698 as moderate risk, 7,020 as high risk, and 4,046 as very high risk.

TABLE 3. MULTIDIMENSIONAL RISK DISTRIBUTION IN THE FULL JOB-POSTING CORPUS

Risk profile	Number of postings	Percentage
Low	4,115	18.81%
Moderate	6,698	30.61%
High	7,020	32.09%
Very High	4,046	18.49%
Total	21,879	100%

Table 3 shows the risk distribution shows that high-risk postings form the largest group (32.09%), followed by moderate-risk cases (30.61%). Low- and very-high-risk postings account for 18.81% and 18.49%, respectively, indicating substantial variation in recruitment-risk exposure across the 21,879-job corpus. Workload feasibility was the dominant dimension for 7,681 postings, followed by

experience rationality for 5,572, role alignment for 4,380, and internal discrepancy for 4,246 cases. The overall risk score was expressed as

$$R_i^{overall} = \frac{R_i^{role} + R_i^{exp} + R_i^{work} + R_i^{disc}}{4}. \quad (12)$$

This is illustrated in Figure 8 below, where a reviewer not only gets to know whether the posting was flagged, but also the recruitment risk that plays the most important part in defining the posting's characteristics. The recruitment risks mentioned above are the indicators of analysis and not illegal or unethical acts.




Fig. 8. Four-Dimensional Recruitment-Risk Profile Incorporating Role Alignment, Experience Rationality, Workload Feasibility, and Internal Discrepancy

3.7 Semantic-Hybrid Representation Experiment

Table 4 shows that semantic information captures predictive structure not completely represented by the 14 manually engineered variables. The hybrid representation produced the highest accuracy, precision, F1-score, and MCC, whereas the semantic-only representation achieved slightly higher recall and discrimination metrics. The results therefore indicate complementarity between interpretable structured features and latent textual information.

TABLE 4. COMPARISON OF HANDCRAFTED, SEMANTIC, AND HYBRID REPRESENTATIONS

Representation	Accuracy	Precision	Recall	F1	MCC	ROC-AUC	PR-AUC
Handcrafted features	80.86 %	94.40 %	79.56 %	86.35 %	0.574	0.901	0.968
Semantic TF-IDF + SVD	94.26 %	96.82 %	95.60 %	96.20 %	0.845	0.9840	0.9946
Hybrid semantic + handcrafted	94.50 %	98.36 %	94.34 %	96.31 %	0.858	0.9836	0.9944

The hybrid experiment is treated as a representation-extension experiment rather than being used to retrospectively replace the primary calibrated LightGBM analysis.

3.8 Robustness and Potential Bias-Trigger Analysis

Table 5 indicates stable predictive performance across 2,000 bootstrap resamples, with relatively narrow confidence intervals for F1, ROC-AUC, and PR-AUC. The wider MCC interval reflects greater uncertainty associated with class imbalance and false-positive variation in the held-out test sample.

TABLE 5. BOOTSTRAP ROBUSTNESS OF THE CORE MODEL

Metric	Observed value	Approx. 95% bootstrap interval
F1-score	0.919	0.895–0.939
MCC	0.621	0.528–0.707
ROC-AUC	0.939	0.916–0.961
PR-AUC	0.980	0.970–0.988

Potential bias-trigger screening identified 17 postings, including 12 with age-related wording and five with exclusionary personality expressions. These signals were

treated only as contextual review prompts; they do not establish demographic bias or discrimination without protected-group attributes and observed hiring outcomes.

IV. CONCLUSION

This study provides an explainable and confidence-adaptive AI framework for recruitment-risk screening, which incorporates predictive modelling, calibrated uncertainty, SHAP-based interpretation, multidimensional risk profiling and Human-in-the-Loop oversight. The calibrated LightGBM model achieved an accuracy of 87.08%, recall of 97.16%, F1-score of 91.94%, ROC-AUC of 0.939 and PR-AUC of 0.980 on the held-out test set. The confidence-based routing required substantive review of only 44.98% of the cases while capturing 96.30% of the observed model errors, thus demonstrating the value of selective human intervention. Furthermore, the improved performance was observed when using a semantic-hybrid representation, achieving an F1-score of 96.31% and an MCC of 0.858. Nevertheless, the current solution is still a decision-support system rather than a fully automated hiring tool due to the usage of weak supervision, single dataset and retrospective Human-in-the-Loop analysis. The future work should focus on the external validation, application of expert-labelled data, prospective recruiter evaluation and applicant-level fairness analysis.

REFERENCES

- [1] A. Mahade, A. Elmahi, K. M. Alomari, and A. A. Abdalla, "Leveraging AI-driven insights to enhance sustainable human resource management performance: Moderated mediation model: Evidence from UAE higher education," *Discover Sustainability*, vol. 6, no. 1, 2025, doi: 10.1007/s43621-025-01114-y.
- [2] S. Yaseen, M. Rao, S. A. Shahani, and A. Javed, "The impact of explainable AI-based interviews on organizational attractiveness: The mediating roles of perceived organizational support and innovativeness," *International Journal of Advanced Social Studies*, vol. 6, no. 1, pp. 117–132, 2026, doi: 10.70843/ijass.2026.06110.
- [3] M. Sandeep, V. Lavanya, and J. Balakrishnan, "Leveraging AI in recruitment: Enhancing intellectual capital through resource-based view and dynamic capability framework," *Journal of Intellectual Capital*, vol. 26, no. 2, pp. 404–425, 2025, doi: 10.1108/JIC-05-2024-0155.
- [4] B. Gül and A. Soyer, "AI-based recruitment: An integrative framework for human resources professionals' adoption," *Systems*, vol. 14, no. 6, Art. no. 713, 2026, doi: 10.3390/systems14060713.
- [5] A. A. Zahrani, "AI in HRM: A unified model of AI capabilities and process efficiency for enhanced HRM effectiveness," *Research Square*, 2025, doi: 10.21203/rs.3.rs-7685386/v1.
- [6] N. A. O. Opong, K. K. Yeboah, J. M. Opong, and E. K. Teye, "Artificial intelligence and people analytics in human resource decision-making: Opportunities, challenges, and ethical imperatives," *IRASS Journal of Multidisciplinary Studies*, 2026, doi: 10.67564/irassjms.v3.i7.0180.
- [7] C. Bura, P. Naayini, and S. Kamatala, "Transforming entrepreneurial human resource management: AI-driven solutions and future directions," in *Advances in Intelligent Systems Research*, pp. 577–592, 2025, doi: 10.2991/978-94-6463-787-8_45.
- [8] L. Hofeditz, S. Clausen, A. Rieß, M. Mirbabaie, and S. Stieglitz, "Applying XAI to an AI-based system for candidate management to mitigate bias and discrimination in hiring," *Electronic Markets*, vol. 32, no. 4, pp. 2207–2233, 2022, doi: 10.1007/s12525-022-00600-9.
- [9] W. Rodgers, J. M. Murray, A. Stefanidis, W. Y. Degbey, and S. Y. Tarba, "An artificial intelligence algorithmic approach to ethical decision-making in human resource management processes," *Human Resource Management Review*, vol. 33, no. 1, Art. no. 100925, 2022, doi: 10.1016/j.hrmr.2022.100925.
- [10] G. Panda, M. K. Dash, A. Samadhiya, A. Kumar, and E. Mulat-Weldemeskel, "Artificial intelligence as an enabler for achieving human resource resiliency: Past literature, present debate and future research directions," *International Journal of Industrial Engineering and Operations Management*, vol. 6, no. 4, pp. 326–347, 2023, doi: 10.1108/IJIEOM-05-2023-0047.
- [11] A. Benabou and F. Touhami, "Integration of artificial intelligence in human resource management: Analyzing opportunities, challenges, and human-AI collaboration," *Management and Accounting Review*, vol. 24, no. 1, 2025, doi: 10.24191/mar.v24i01-15.
- [12] T. Radouche, "How AI enabled human resource management drives sustainable hybrid performance through AI-centric green human resource management in Saudi Arabian private-sector firms," *Discover Sustainability*, 2026, doi: 10.1007/s43621-026-04475-0.
- [13] P. N. Tinguely, J. Lee, and V. F. He, "Designing human resource management systems in the age of AI," *Journal of Organization Design*, vol. 12, no. 4, pp. 263–269, 2023, doi: 10.1007/s41469-023-00153-x.
- [14] F. Aziz, F. Muzaffar, S. Shahid, H. S. Ahmed, and S. M. Iqbal, "The role of artificial intelligence in driving ROI through synergized HR, marketing, and financial decision-making," *Inverge Journal of Social Sciences*, vol. 4, no. 3, pp. 129–142, 2025, doi: 10.63544/ijss.v4i3.153.
- [15] A. Bujold, I. Roberge-Maltais, X. Parent-Rochelleau, J. Boasen, S. Sénécal, and P. Léger, "Responsible artificial intelligence in human resources management: A review of the empirical literature," *AI and Ethics*, vol. 4, no. 4, pp. 1185–1200, 2023, doi: 10.1007/s43681-023-00325-1.