systems

MDPI

Article

Humanizing ATS-Based Recruitment Using LLMs and Human-in-the-Loop Oversight

Valdo V. Mpinga [1,*] and António Miguel Rosado da Cruz [2,*]

[1] ESTG, Instituto Politécnico de Viana do Castelo, 4900-348 Viana do Castelo, Portugal
[2] ADiT-Lab, Instituto Politécnico de Viana do Castelo, 4900-348 Viana do Castelo, Portugal
* Correspondence: valdom@ipvc.pt (V.V.M.); miguel.cruz@estg.ipvc.pt (A.M.R.d.C.)

Abstract

Application Tracking Systems (ATSs) have evolved significantly since their inception in 1996, transitioning from simple resumérepositories to AI-driven tools with advanced capabilities. While these developments have improved recruitment efficiency, they have also raised important ethical, organizational, and human-rights-related concerns. Bias in machine learning (ML) training data, opaque decision criteria, and excessive reliance on automated judgment may contribute to unfair treatment, reduced transparency, and limited human oversight in hiring processes. This study addresses these challenges by proposing a human-centered approach to ATS-supported recruitment based on a set of Humanization Services. Using a Design Science Research approach, three main artifacts were developed: a Job Requirements Validation Module, a Bias Trigger Removal Module, and a blockchain-supported dual-authorization mechanism for vacancy approval, which requires digital signatures from qualified professionals to approve job postings, ensuring that there are humans that assume responsibility. These components are intended to improve job posting quality, reduce bias-conducive information in applicant data, and strengthen accountability in recruitment workflows. The evaluation provides initial empirical support for the operational feasibility of the proposed approach under the tested conditions. The study therefore contributes a practical and theoretically grounded step toward more transparent, accountable, and human-centered AI-supported recruitment.

Keywords: application tracking system (ATS); artificial intelligence (AI); business process model and notation (BPMN); human resources management (HRM); large language model (LLM); machine learning (ML)

1. Introduction

Application Tracking Systems (ATSs) are a type of software used by companies, typically by management and human resources (HR) teams to efficiently manage job openings, respective job advertisements, and the applications for those openings. In 1996, the first ATSs appeared, with features for storing applicants' resumés, keyword filtering and searching, advertising job openings, and candidate tracking [1]. Years later, with the advancement of technology, this type of software evolved and was able to do much more. The rapid development of Artificial Intelligence (AI), especially Machine Learning (ML), and its integration into ATSs have enabled an increase in productivity in HR management teams. Due to the AI/ML integration into ATSs and the way ML algorithms are trained, several problems emerged. Many of these problems have to do with the training datasets used, which may include data with racial, sexual, and age biases. These have also created a

Check for updates

Received: 14 March 2026
Revised: 14 April 2026
Accepted: 17 April 2026
Published: 22 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

certain dependence and accommodation in HR management teams due to the practice of recruitment based mainly on algorithmic management.

Unfortunately, these problems left the technical scope and entered the ethical human realm. The problem is not the use of these tools but how the systems themselves are made and designed as they violate human dignity through unfair algorithmic processes [2] and violate basic human rights, such as the right to work, equality and non-discrimination, privacy, right of expression and right of free association [3–5].

Despite progress in AI-driven recruitment, the conclusions from the Related Work review section show us that key issues remain unresolved. These include the lack of validated ethical frameworks, unclear human–AI decision balance, inconsistent fairness definitions, and weak regulatory standards.

This research work focuses on developing empirically validated, human-centered ATS pluggable components that, without compromising the recruitment efficiency enabled by ATSs, uphold fairness, transparency, and respect for human dignity. For this purpose, a set of independent services for evaluating and humanizing job advertisements is proposed. One of these services aims to identify misaligned requirements, unreasonable expectations and unrealistic workloads. Another service detects bias triggers in recruitment data, such as gender, age, race and other personal characteristics, which can lead to discriminatory practices. By applying AI and data validation techniques, the service aims to remove biased data that do not contribute to the assessment of an individual's true professional potential. Furthermore, it ensures that job advertisements are rational, fair, and aligned with the actual requirements for the role, thereby promoting a more inclusive and equitable hiring process.

This research work also proposes a solution to address the growing dependence on algorithmic recruitment, encouraging human involvement in the final selection process. This approach seeks to strike a balance between AI efficiency and the critical need for human judgment to make fair and unbiased hiring decisions.

Building upon this foundation, the rest of this article is structured as follows. Section 2 discusses various studies addressing algorithmic bias, the lack of transparency, violations of human dignity and rights, and the importance of ethical considerations and regulatory frameworks in AI recruitment. It also explores technological advancements like blockchain for enhancing trust and accuracy, the perceptions and acceptance of AI in HRM, methods for improving resumé parsing, and strategies for mitigating bias-conducive factors such as age, gender, and race. The methodology, detailed in Section 3, employed a Design Science Research approach to create and evaluate a humanization service application. Section 4 describes the design and implementation of the humanization service application. The validation and results, presented in Section 5, highlight findings from the analysis of a large dataset of job postings, revealing prevalent issues. This section also discusses the evaluation of different language models, the functional verification of the bias triggers removal service, the robust testing of the blockchain validation system, and the demonstration of the system's capabilities through the public web interface. Concluding the article, Section 6 summarizes the key findings and contributions of the research, emphasizing the development and validation of the humanization service application as a means to address bias and dehumanization in AI-driven recruitment, ultimately aiming for fairer and more transparent hiring practices.

2. Related Work

The problems with recruitment relying on ATSs have caused awareness among several researchers, culminating in several papers discussing the violations and possible approaches to address the unfairness and human rights violations caused by excessive

reliance on ML algorithms. The problem is not at all about the use of tools and algorithms as these algorithms have been developed to improve the recruitment process for both recruiters and applicants. The problem is mostly how the system is used and the bias on the dataset used for training the ML models.

Research on AI-driven recruitment has progressively shifted from a narrow focus on classifier bias to a broader view of recruitment as a socio-technical pipeline. Recent surveys argue that unfairness in ATS-based hiring can emerge not only in candidate screening but also in job-ad design, recommender-system exposure, resumé parsing, ranking, and downstream human review [6–9]. This broader perspective helps explain why the central problem is not the mere use of tools and algorithms but rather how these tools are designed, trained, deployed, and governed, including the quality and representativeness of the data used to train them.

This view is consistent with earlier work emphasizing the ethical and legal implications of AI-driven recruitment. With the goal of preventing discrimination, and motivated by the lack of regulation in AI-driven recruitment, the authors in [10] developed a Multi-Agent System architecture to support ethical and legal auditing of AI-driven recruitment, with particular attention to video-interview analysis. Related works have further shown that excessive reliance on algorithmic assessment can undermine human dignity and infringe fundamental rights, including equality and non-discrimination, privacy, freedom of expression, freedom of association, and the right to work [2,11]. In this sense, the main concern is not only technical bias but also the reduction of candidates to narrow and fixed machine-readable profiles.

Another recurring theme in the literature is the need to balance automation and human judgment. When recruitment is driven solely by algorithmic management, applicants tend to perceive the process as less fair [12]. Paramita et al. showed that both transactional and relational aspects of AI-mediated recruitment matter for the perceived quality of recruitment services [13]. More recent experimental work reinforces this point: applicants generally perceive AI-only resumé screening as less fair than human-led or balanced human–AI screening, especially in rejection scenarios [14,15]. These findings suggest that meaningful human oversight is not merely a regulatory safeguard but also an important determinant of perceived procedural justice.

Bias remains one of the most intensively studied problems in algorithmic hiring. In [6], Fabris et al. examine bias-conducive factors in algorithmic hiring and highlight the importance of group fairness, granularity, and interpretability when evaluating fairness measures. Their analysis confirms that characteristics such as age, disability, gender, religion or belief, racial or ethnic origin, and sexual orientation can all become sources of unfairness in recruitment systems [6].

More recent work extends this concern to contemporary language-model-based systems. Wilson and Caliskan show that language-model retrieval for resumé screening can reproduce gender, race, and intersectional bias, while Seshadri et al. show that LLM-based hiring systems can display allocational unfairness in both resumé summarization and ranking tasks [16,17]. These findings indicate that moving from keyword-based ATS pipelines to embedding-based or generative systems does not eliminate bias. Instead, it shifts bias into newer layers that also require auditing.

At the same time, the literature does not support a purely anti-technology conclusion. In fact, adoption of AI is not the root cause of the problem. It is often how some organizations approach and use some of the AI tools. In [18], Prasad et al. investigated the impact of generative AI tools on HRM practices, organizational commitment, employee engagement, and performance, exploring the mediating role of trust in the relationship between user perception and organizational commitment. They concluded that the adoption of AI tools can

enhance organizational commitment, employee engagement, and employee performance, providing valuable information on the adoption of AI tools in the workplace [18].

Tsiskaridze et al. likewise note growing interest in AI for HRM because of its potential to mitigate some forms of individual human bias and support more consistent decision making [19]. However, these benefits depend on transparency, fairness, explainability, and appropriate institutional safeguards. Similarly, Aleisa et al. propose that integrating AI with blockchain technology may improve transparency and trust in recruitment workflows [20].

At the operational level, accurate extraction of candidate information also matters: Kinger et al. propose a refined ATS pipeline with improved resumé parsing performance, reporting 96.2% accuracy in parsing and thereby significantly improving the efficiency of the candidates selection process [21].

A closely related but less mature line of research moves upstream from candidate screening to the integrity of job postings themselves. Recent work shows that job advertisements should not be treated as neutral inputs to ATSs. Frissen et al. developed machine-learning methods to identify biased and discriminatory language in job advertisements, including masculine-coded, feminine-coded, exclusive, LGBTQ-coded, demographic, and racial language [22]. This is important because bias introduced at the advertisement stage can affect both who applies and who is later judged by the ATS. More broadly, recent recruitment research increasingly treats job ads, recommendation systems, and screening models as interconnected parts of the same fairness problem [6,8,9].

Recent NLP work also provides foundations for validating the integrity of job ads in a more structured way. Senger et al. survey computational methods for extracting and classifying skills from job postings, and Zhang et al. introduce the *SkillSpan* benchmark for extracting hard and soft skills from English job ads [23,24]. These studies make it possible to transform a free-text posting into structured claims about tasks, skills, and qualifications. Building on that idea, Urbano et al. propose methods for inconsistency detection in job postings, showing that NLP and rule-based approaches can identify contradictory or ambiguous requirements before publication [25]. Although this line of work is still emerging, it is especially relevant for detecting unrealistic requirements, contradictory expectations, and misleading qualification bundles.

Related evidence also suggests that job posting quality and integrity depend on the completeness and informativeness of the content provided to applicants. Arnold et al. show that pay-transparency requirements materially increase the prevalence of salary information in postings and affect labor-market outcomes [26]. Audoly et al. further show that both pay and non-pay content in job ads are informative about employer attractiveness, indicating that omissions in pay, contract duration, flexibility, and other working conditions are not trivial formatting choices but important informational signals [27]. In parallel, Naudé et al. demonstrate that machine-learning methods can distinguish different forms of fraudulent job advertisements, suggesting that fake-job detection should also be seen as part of job posting integrity assessment [28].

Finally, user acceptance of AI in recruitment remains mixed. Nastase et al. report that most variables in their model positively influenced employees' intention to accept and use AI in recruitment and selection, although non-discrimination and the role of AI in these processes had only limited influence [29]. This finding is noteworthy because it suggests that organizational acceptance of AI does not automatically imply adequate attention to fairness risks.

In summary, the literature has made substantial progress in identifying the ethical, legal, and technical challenges of AI-driven recruitment and ATSs, but several issues remain unresolved. First, many ethical-auditing and governance frameworks remain conceptual

and have not yet been validated extensively on real-world recruitment data [7,10,30]. Second, although several studies emphasize the importance of balancing algorithmic and human judgment [11,13–15], the literature still provides limited guidance on how this balance should be operationalized in practice. Third, there is still no consensus on how fairness should be defined, measured, or optimized in hiring, especially when sensitive attributes are partially available, excluded, or deliberately used to support historically disadvantaged groups [6]. Finally, while recent work on biased job-ad language, skill extraction, contradiction detection, pay transparency, and fraudulent postings provides important building blocks [22–28], there is still no widely adopted end-to-end framework for validating the integrity of job postings and systematically detecting unrealistic requirements, contradictory expectations, and exploitative practices.

3. Methodology

This research follows the *Design Science Research* (DSR) methodology [31,32]. DSR is particularly appropriate for this study because it is oriented toward the design, development, and evaluation of artifacts intended to solve relevant real-world problems while simultaneously generating practical and theoretical contributions [31,33]. In the present case, the problem domain is inherently socio-technical: recruitment processes mediated by ATS technologies involve not only technical issues such as model bias, opacity, and unreliable data extraction but also organizational, ethical, and regulatory concerns related to fairness, accountability, and human oversight.

Accordingly, the purpose of this study is not merely to analyze ATS-related problems but to design and evaluate concrete artifacts that help mitigate them. More specifically, the research aims to improve the fairness, transparency, and accountability of AI-supported recruitment by developing a set of Humanization Services that intervene in critical points of the recruitment pipeline. These interventions are intended to ensure that technological advances remain aligned with ethical values, legal requirements, and practical human resource management needs. To structure the study, the six activities of the DSR process proposed by Peffers et al. were followed [31,32] (see Figure 1).

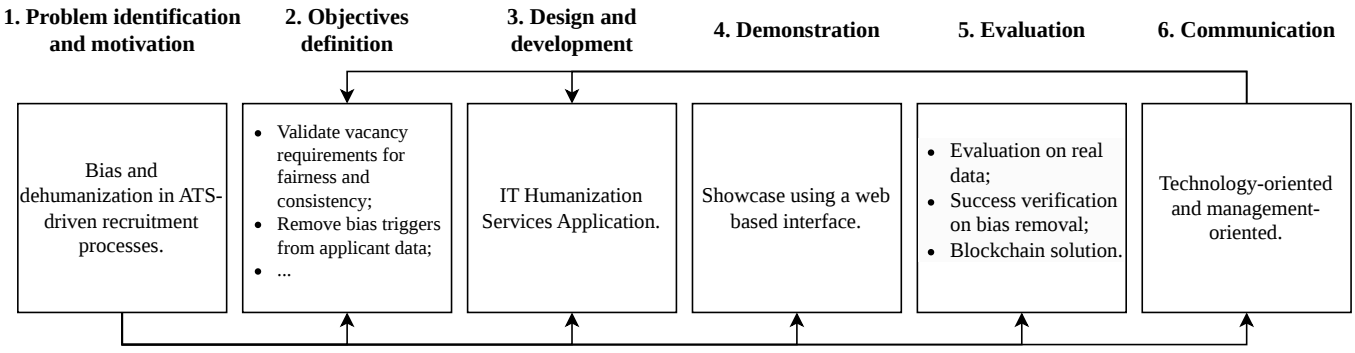


Figure 1. DSR diagram applied to the project's context.

3.1. Methodological and Theoretical Framing

Although DSR provides the methodological structure for the research, the study is also informed by a systems-theoretic and socio-technical perspective. From this viewpoint, ATS-based recruitment is not treated as a single isolated algorithmic decision but as an open organizational system composed of interdependent elements, including job advertisements, applicant data, resumé-parsing tools, ranking models, recruiters, organizational rules, and compliance mechanisms [6–9]. Bias is therefore understood as a systemic phenomenon that may emerge at different stages of the recruitment pipeline, rather than as a defect restricted to one classifier or one dataset.

This perspective is especially relevant for the present study because the proposed artifacts intervene at multiple points in the system. The validation of job postings addresses upstream distortions in the representation of vacancies; the removal of bias triggers from applicant data targets risks in automated screening and ranking; and the blockchain-based validation mechanism introduces accountability and shared human oversight into the publication process. In this sense, the artifacts are designed not only as technical components but as socio-technical mechanisms intended to improve the interaction between automation, human judgment, and organizational governance.

3.2. Problem Identification and Motivation

The first DSR activity consisted of identifying and motivating the problems to be addressed. Based on the literature reviewed in the previous section, this study focuses on three central challenges.

- Algorithmic bias in candidate screening. ATS-based recruitment systems may disadvantage certain groups due to biased training data, proxy variables, or poorly operationalized notions of fairness. This problem is compounded by the lack of consensus on how fairness should be defined, measured, and implemented in AI-assisted hiring, particularly when sensitive attributes such as gender, race, and age are involved [2,6,34].
- Lack of transparency and limited empirical validation. Many AI-based recruitment systems remain opaque to applicants and recruiters, making it difficult to understand how candidates are assessed and reducing trust in the process. This issue is aggravated by the fact that several proposed ethical or auditing solutions remain conceptual and have not been sufficiently validated in real-world contexts [10–12].
- Insufficient balance between human and algorithmic decision-making. Excessive reliance on automated evaluation may weaken ethical judgment, reduce contestability, and compromise human dignity. Preserving a meaningful human role in recruitment is therefore essential to prevent overreliance on machine outputs and ensure that broader organizational and ethical considerations are taken into account [11,13–15].

Addressing these challenges is significant not only from a technical standpoint but also from an organizational and societal one. More equitable and transparent recruitment processes can improve institutional trust, reduce legal and reputational risks, and support the responsible adoption of AI in human resource management. This aligns with the DSR principle of producing innovative and useful solutions to relevant business problems.

3.3. Definition of Objectives

In response to the identified problems, the objective of this research is to design and develop a Humanization Services application capable of introducing fairness, transparency, and accountability into ATS-supported recruitment. The application is intended to achieve three main objectives:

- Validate job opening requirements in order to improve fairness and consistency in job postings. This includes identifying misalignments, unreasonable demands, unrealistic expectations, and internal contradictions in vacancy descriptions.
- Remove bias triggers from applicant data so as to mitigate the risk of discriminatory automated screening. This objective focuses on identifying and suppressing sensitive or bias-conducive information while preserving the structural integrity and usability of the original data.
- Implement a digital signature mechanism for human reviewers to enhance accountability and transparency in the publication of job openings. This is achieved through a

decentralized validation protocol based on blockchain technology, requiring authorization from both HR personnel and subject-matter experts.

These objectives are primarily qualitative in nature as they are concerned with improving the ethical and organizational quality of recruitment processes. At the same time, they are operationalized through the development of concrete artifacts that can be demonstrated and evaluated empirically.

3.4. Design and Development

The design and development phase focused on constructing the proposed Humanization Services as functional artifacts. Three main components were implemented:

- Vacancy Requirement Validation Module. This module was developed using Python 3.13 and FastAPI and incorporates LLaMA-based prompting to evaluate the fairness and consistency of job postings. It analyzes job descriptions and produces structured feedback on misalignments, unreasonable requirements, unrealistic expectations, and internal discrepancies.
- Bias Trigger Removal Module. Also implemented with Python 3.13, FastAPI, and LLaMA-based prompting, this module identifies and removes bias-related or sensitive fields from applicant data while preserving the original format and usability of the input.
- Digital Signature of Relevant Human Actors Module. A decentralized validation protocol was designed using blockchain technology to require dual authorization from HR personnel and subject-matter experts before a vacancy can be published.

Taken together, these components constitute the main artifacts of this research. They represent practical outputs of the DSR process in the form of methods, services, and technical features aimed at humanizing and strengthening ATS-supported recruitment.

3.5. Demonstration

The demonstration phase aimed to show that the developed artifacts can operate in a realistic setting and address the identified problems in practice. The Humanization Services application was demonstrated through the following activities:

- Validating a set of job postings to illustrate the system's ability to detect inconsistencies, unfair requirements, and potential bias triggers;
- Processing applicant data to demonstrate the removal of bias-conducive information while preserving data integrity;
- Simulating the blockchain-based validation workflow to confirm the correct implementation of access control, authorization, and signature verification;
- Presenting a web interface that integrates job analysis, bias-reduced candidate visualization, and blockchain validation (available at https://joblimpo.valdompinga.com/ (accessed on 15 January 2026)).

These demonstrations provide evidence that the proposed artifacts are not merely conceptual but operational and suitable for practical use.

3.6. Evaluation

The evaluation phase assessed the extent to which the proposed artifacts achieved the objectives defined earlier. Evaluation focused on the effectiveness of the Humanization Services in improving fairness, transparency, and workflow accountability.

- The Vacancy Requirement Validation Module was evaluated on a dataset of 21,701 job postings, measuring the proportion of postings that presented misalignments, unreasonable demands, unrealistic expectations, or internal discrepancies.

- The Bias Trigger Removal Module was evaluated by verifying the successful identification and removal of bias-related fields from applicant data while maintaining the consistency of the original structure.
- The Digital Signature Module was evaluated through tests of access control, signature verification, and workflow integrity in the blockchain-based validation process.

The evaluation compared the observed results with the objectives established in the design phase and employed analysis procedures appropriate to each artifact. In line with DSR, the purpose of the evaluation was not only to verify technical functionality but also to assess whether the artifacts meaningfully contribute to addressing the identified socio-technical problems.

3.7. Communication

The final DSR activity consists of communicating the research process and results to relevant audiences. In this study, communication is carried out through the present article, which reports the identified problems, the conceptual framing, the design of the Humanization Services application, and the results of its demonstration and evaluation.

Communication is intentionally addressed to both technology-oriented and management-oriented audiences. For technology-oriented readers, such as AI developers and software engineers, the article provides details on the architecture, implementation technologies, prompting strategy, and blockchain-based validation mechanism. For management-oriented readers, such as HR professionals and organizational decision-makers, the focus is placed on how the proposed artifacts can improve fairness, increase transparency, reduce legal and reputational risks, and support more trustworthy recruitment processes.

4. Design and Implementation

The core objective of this study is to improve fairness and enhance human involvement in key recruitment processes. This is achieved by pre-processing job requirements before they are published, by addressing biased applicant's data before they are analyzed by the ATS, and through the involvement and accountability of relevant human actors. This is made possible through a Humanization Service API whose main services are revealed in Table 1. Also, a Humanization Service Application has been developed, incorporating the necessary features to allow the direct use of these enhancements.

1. Validation of vacancy requirements to be published—Contemporary analysis of the employment landscape reveals a prevalent issue: the dissemination of job vacancies characterized by incongruous and often unattainable prerequisites. These discrepancies range from entry-level positions, stipulating multi-year experience levels, to roles demanding expertise exceeding the temporal existence of the relevant industry or technology. To address these systemic inconsistencies, an intelligent automation framework was developed for rigorous evaluation of job vacancy postings. This framework undertakes a multifaceted assessment to identify potential contradictions between designated role titles and articulated requirements. Beyond the detection of unrealistic experience demands, the system was engineered to scrutinize job descriptions for a broader spectrum of potential issues. This includes the identification of misaligned skill sets, evaluation of workload feasibility within the scope of a single position, and detection of any internal inconsistencies within the vacancy description itself. Furthermore, the analytical capabilities extend to the formulation of actionable recommendations for rectifying identified issues, such as suggested adjustments to the role title, modifications to specific requirements, and revisions to experience-level expectations. Critically, the framework incorporates a module dedicated to the identification of potential violations of ethical and human-centered employment practices,

ensuring a more equitable and transparent recruitment process. The output of this automated validation process encompasses a comprehensive evaluation, including a detailed breakdown of identified discrepancies, a set of targeted recommendations for improvement, and an overall assessment of the vacancy's compliance with the established criteria.

Table 1. API services for validating job requirements, removing bias triggers and digital dual-signature of job requirements by relevant human actors.

API Service	Short Description	Input Parameters	Output
Validation of vacancy requirements	Reports contradictions between designated role titles and defined requirements, among other potential issues (e.g., unrealistic experience demands, misaligned skill sets, workload feasibility, inconsistencies, etc.) within the vacancy description.	Specifications for a given job opening (in JSON format)	JSON object with identified discrepancies, recommendations for improvement, and an overall assessment of the vacancy's compliance with the established criteria.
Mitigation of bias triggers	Identification and removal of identified bias-inducing factors within the recruitment's or candidate's data.	Applicant data or resumé (in JSON, XML or plain text format)	Applicant data or resumé, in the same format as the input data (JSON, XML or plain text format), devoid of bias-inducing information.
Digital signature of relevant human actors	Blockchain-based protocol for ensuring validation and responsibilization of both HR personnel and a field expert.	JSON object with job title, job description, job type, HR signature, field manager signature	Notification of job approval/rejection, invalid HR or field manager signature, or missing field manager for the required job type.

2. Mitigation of Bias Triggers—A significant concern in contemporary hiring practices pertains to the potential for discriminatory biases arising from the collection and utilization of sensitive personal information. Attributes such as age, gender identity, sexual orientation, and racial or ethnic background have historically served as triggers for prejudiced decision-making. To counteract these inequitable scenarios, a methodology focused on the identification and subsequent reduction in bias triggers within recruitment data has been developed. This proactive approach aimed to facilitate the development and refinement of recruitment models that operate with enhanced fairness and impartiality. After the job opening requirement analysis, this "Mitigation of Bias Triggers" dedicated mechanism was implemented to process the candidate's data. This mechanism is specifically designed to ingest input data, meticulously preserve the original structural format, and systematically eliminate attributes recognized as potential sources of bias. These attributes include but are not limited to name, age, gender, sexual orientation, race or ethnicity, religious affiliation, disability status, and marital or parental status, thereby promoting a more equitable evaluation of candidate qualifications.
3. Digital signature of the relevant human actors—A significant impediment to fair and efficient talent acquisition lies in the publication of inconsistent, inflated, or unrealistic job requirements. Such deficiencies may discourage suitable applicants, distort downstream screening criteria, and contribute to unfair hiring outcomes. To mitigate this risk, the proposed system incorporates a validation mechanism in which a human subject-matter expert with relevant domain knowledge reviews the vacancy requirements after the automated humanization process. This additional stage of scrutiny helps ensure that the requirements are technically sound, proportionate to

the role, and aligned with actual organizational needs. To operationalize this step, a graphical user interface is provided through which at least one qualified actor, such as a project manager or technical specialist, must explicitly approve the vacancy prior to publication. The integrity and traceability of this approval workflow are supported by a blockchain-based distributed ledger, which provides a tamper-evident and non-repudiable record of the approval events. Thus, blockchain is employed not to detect bias directly but to strengthen accountability, enforce human oversight, and preserve an auditable history of vacancy authorization.

4.1. Architecture

The proposed system architecture illustrated in Figure 2 delineates the implementation of the aforementioned functionalities as a modular service, thereby facilitating seamless integration into existing recruitment ecosystems. On the left-hand side of the diagram, an entity or corporation utilizing an ATS is depicted as the primary consumer of the ATS Humanization Service. This service is designed as a centralized module offering three service endpoints: a comprehensive analysis of vacancy requirements; the systematic removal of bias-related data from the applicants' information; and a digital signature for relevant human actors to subscribe and validate a job specification (refer to Table 1). Interaction is initiated when the ATS transmits either a vacancy description or applicant's resumé to the Humanization Service. Within this service, a dedicated Humanization Service component orchestrates initial processing. Subsequently, a Data Processing module undertakes the core analytical tasks, potentially leveraging an LLM for sophisticated text analysis and pattern recognition. The output of this process, humanized data, which may encompass insights derived from vacancy analysis or bias-redacted candidate information, is then relayed back to the originating ATS. When posting a job vacancy another service endpoint allows for human validation and responsibilization for its contents, by both HR personnel and a field expert. This service-oriented architecture promotes modularity and reusability, allowing for straightforward incorporation of advanced humanization capabilities into diverse ATS platforms, by calling the service's API.

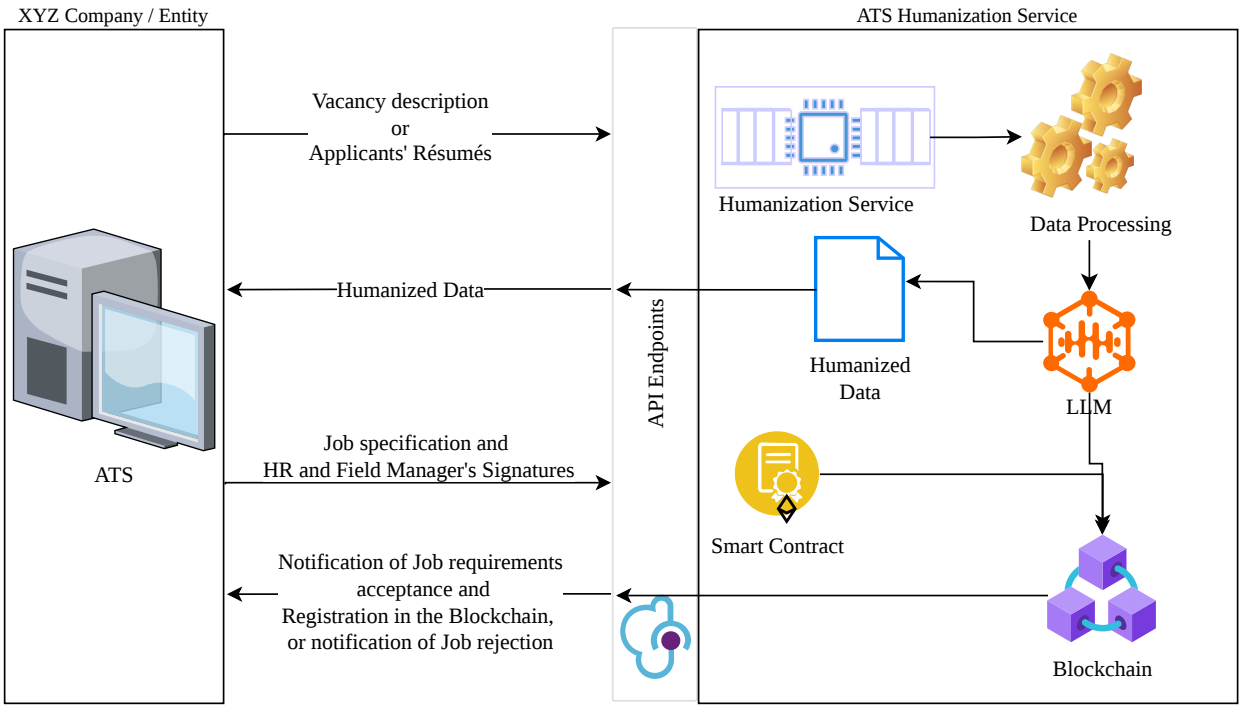


Figure 2. Solution architecture.

The main features implemented to enhance humanization in this system are the validation of vacancy requirements before publication, and the identification of bias triggers. To address these challenges efficiently, artificial intelligence has been integrated into the process, specifically using a Large Language Model (LLM) called LLaMA (large language model Meta AI). Developed by Meta-AI, LLaMA is a natural language processing (NLP) model designed to advance research in generative AI [35]. It is a transformer-based model trained on an extensive and diverse corpus of text, offering high performance with fewer parameters than other large-scale models, such as GPT-4, while achieving competitive results across various NLP benchmarks [35].

Developers interact with LLMs primarily through prompts, which are textual inputs or instructions that guide the model's response. These prompts can vary in complexity from simple queries to detailed commands, and the level of specificity directly affects the quality of the output. When processing a prompt, the input is tokenized, a process in which the text is broken down into smaller units or tokens, enabling the model to interpret and generate a response accurately.

The rise of open-source LLMs has opened up numerous opportunities as they allow developers to accelerate workflows. Tasks that require coding small, repetitive features can now be simplified with the right prompt, which can produce a functional output efficiently. However, a significant limitation of LLMs is their consistency. Owing to their probabilistic nature, these models generate responses word by word based on the likelihood of each word following the previous one [36]. Consequently, the output can vary even when the same prompt is used multiple times.

Fortunately, this inconsistency can often be mitigated in scenarios that require predictable behavior. By crafting well-structured prompts that include specific rules or formats, a model is more likely to produce consistent outputs. This makes LLMs highly useful for smaller, well-defined tasks and systems, where consistency is essential. When the prompt sets clear expectations, the model can reliably generate responses in the desired format, thus making it a practical tool for many applications. Section 5.3 analyzes the performance, accuracy, and result consistency of several open-source LLMs in job requirement analysis.

4.2. Interactive Web-Based Frontend

To effectively illustrate the practical applicability and facilitate user interaction with the proposed solution, a publicly accessible web interface was developed and is accessible at the URL https://joblimpo.valdompinga.com/ (accessed on 15 January 2026). This frontend serves as a demonstration platform and the primary point of interaction, allowing users to input queries, configure parameters, and visualize the outputs generated by the underlying system.

The user interface of the system was implemented as an interactive website designed to provide intuitive access to its core functionalities. The web demonstration interface is structured around four principal pages, each serving a distinct purpose in facilitating the validation of vacancy requirements and the mitigation of bias in the candidate data. The Open-Source Repository for the project's UI is https://github.com/ValdoMpinga/clean-job-UI (accessed on 15 January 2026). And the backend modules can be found at https://github.com/ValdoMpinga/recruitment-humanization-service/ (accessed on 15 January 2026).

The Landing Page serves as the initial point of contact for the users (https://joblimpo.valdompinga.com/ (accessed on 15 January 2026)). This page provides an overview of the key features of the application, highlighting its capabilities in enhancing fairness and objectivity in the recruitment processes.

The Job Opening Requirements' Validation Page, accessible at https://joblimpo.valdompinga.com/requirements (accessed on 15 January 2026), allows any user to analyze a job description against predefined ethical and practical criteria. Users can input the text description or requirements of a job and receive an evaluation of its compliance, identifying potential issues related to role alignment, experience rationality, workload feasibility, and internal discrepancies.

The Candidate Data Bias Removal Page, found at https://joblimpo.valdompinga.com/candidate (accessed on 15 January 2026), provides a feature for users to process candidate-like data with the aim of removing or anonymizing information that could potentially trigger biases in automated applicant tracking systems. This functionality supports a more objective assessment of candidate qualifications.

Finally, the Ethical Validation Page, located at https://joblimpo.valdompinga.com/validation (accessed on 15 January 2026), provides information regarding ethical considerations and the implementation of a blockchain-based solution for handling the validation of code signatures on platforms such as GitHub. This page emphasizes the system's commitment to transparency and integrity in its operations.

These four main pages collectively offer a comprehensive and user-friendly interface for interacting with the system's functionalities, from understanding its core features to actively utilizing its tools for vacancy analysis and bias mitigation. The website adheres to responsive design principles, ensuring accessibility across various screen sizes.

4.3. Validation of Job Requirements for Publication

This feature receives a job requirements description typically generated by HR personnel or an Applicant Tracking System (ATS). These posts often lack alignment with realistic human-centered expectations. The endpoint evaluates the vacancy to ensure that it meets fairness and rationality standards consistently.

To achieve this, a carefully crafted prompt was developed and refined to produce consistent outputs in a standardized format, specifically in a JavaScript Object Notation (JSON) structure. For building both prompts in this study, techniques from [37] have been applied in an iterative manner. The final version of the prompt used for this endpoint is as follows:

> *"You are a Job Requirement Validator. Your task is to evaluate and assess job posts and evaluate them for fairness, rationality, and alignment with the title of a role. Specifically, you should:*
>
> - ***Role Alignment:*** *Check if the listed job requirements are relevant to the job title.*
> - ***Experience Rationality:*** *Ensure the experience requirements are reasonable for the role level.*
> - ***Workload Feasibility:*** *Assess whether the listed responsibilities and requirements are realistic for a single role.*
> - ***Discrepancy Check:*** *Identify inconsistencies or contradictions.*
> - ***Human-Centered Feedback:*** *Highlight any exploitative practices.*
>
> *Evaluate the following job: Job Posting:* `{ATS VACANCY GOES HERE}`
>
> *Return the analysis in the following structured format:*
>
> - ***Role Title:*** *Job Title*
> - ***Metrics:***
>   - ***Role Alignment:***
>   - *Status: Pass/Fail*
>   - *Issues: List of misaligned requirements*

  - ***Experience Rationality:***
    - *Status: Pass/Fail*
    - *Issues: Details about unreasonable experience requirements*
  - ***Workload Feasibility:***
    - *Status: Pass/Fail*
    - *Issues: Details about unrealistic workloads*
  - ***Discrepancies:***
    - *Status: Pass/Fail*
    - *Issues: Details about discrepancies*
- ***Recommendations:***
  - *Role Title Adjustment: Suggested new title if necessary*
  - *Requirement Changes: Suggested changes to requirements*
  - *Experience Changes: Suggested changes to experience requirements*
  - *Other Recommendations: Additional advice or changes*
- ***Violations:***
  - *Human Rights: List of detected violations, if any*
- ***Summary:***
  - *Overall Feedback: Summary of the evaluation*
  - *Compliance Score: Percentage of compliance based on metrics (0–100)*
    PS: Just return the JSON, only!"

This ensures that the model consistently outputs responses in the desired format, making the validation process reliable.

4.4. Revealing Bias Triggers

The goal of this feature is to eliminate fields that could introduce bias during decision-making, such as name, age, sex, and other personal attributes. The endpoint receives structured input data (e.g., JSON, Extensible Mark Language (XML), or plain text) and returns the same data format, but with the bias-triggering fields removed.

"You are an AI tool designed to clean structured data by removing bias-related fields. Bias-related fields include, but are not limited to:

- *Name*
- *Age*
- *Gender*
- *Sexual orientation*
- *Race or ethnicity*
- *Religion*
- *Disability status*
- *Marital or parental status (e.g., "marital_status", "children")*
- *Any photos or physical descriptions.*

Output Rules:

1. *Return the cleaned data in* ***exactly the same format*** *as the input (JSON, XML, or plain text).*
2. *Do* ***not*** *include any explanations, code, examples, comments, or extra text—only the cleaned data.*
3. *Do not format the output with code fences (e.g., "'') or any surrounding markdown or comments.*

4. *If the input is JSON, return valid JSON.*
5. *If the input is XML, return valid XML.*
6. *If the input is plain text, return the cleaned plain text.*
7. *Don't output keys with blank value because of the removal, just remove both keys and value if it has bias data.*
8. *Languages spoken are not bias.*

Input:
{APPLICANT DATA GOES HERE}

Output:
(Return the cleaned input data format strictly as specified.)"

This prompt ensures that the model consistently returns the input data in the same format but is devoid of any bias-related information.

Using these two endpoints, the system addresses critical challenges: ensuring fairness and human-centered evaluation in job postings and removing bias from applicant data. Both implementations demonstrated how carefully designed prompts can enable Large Language Models to perform specific tasks effectively and reliably.

4.5. Digital Signature of Relevant Human Actors

Traditional methodologies for creating job posts frequently suffer from the inclusion of inflated, unrealistic, or misaligned requirements, often stemming from a lack of sufficient domain-specific knowledge during the drafting phase [38,39]. These shortcomings may discourage qualified candidates, misrepresent the actual needs of the role, and introduce unfair criteria into subsequent ATS-based screening stages.

To mitigate this critical issue, we advocate that the final decision and responsibility for a job posting should rest with a human, rather than to automated mechanisms alone. For ensuring this, we propose a decentralized validation protocol based on the necessity of dual cryptographic authorization from both HR personnel and a subject-matter expert, prior to the formal publication of any job vacancy. This is similar to a decentralized signature from the HR manager and a domain-specific expert, ensuring that both are accountable for the content of the job offer. As a result, accountability for the published content becomes distributed, explicit, and verifiable, reducing the likelihood that problematic requirements are approved without appropriate review.

Importantly, the blockchain layer is not intended to improve the semantic quality of the vacancy analysis itself. Instead, it is used to preserve a non-repudiable and tamper-evident record of the human approvals required for publication.

4.5.1. Decentralized System Architecture

The proposed system strategically leverages the capabilities of Ethereum smart contracts to rigorously enforce several key operational parameters, the details of which are listed in Table 2.

- Role-Based Access Control: Implementation of distinct permission frameworks tailored for HR managers and field-specific experts.
- Multi-Signature Validation: Utilization of the Elliptic Curve Digital Signature Algorithm (ECDSA) to ensure robust multi-signature verification [40,41].
- Immutable Approval Records: Secure and transparent recording of all approval processes through the immutable state transitions inherent to the blockchain.

Table 2. Smart contract characteristics.

Characteristic	Details
Digital signature scheme	ECDSA (secp256k1 curve) via `ecrecover`
Smart contract functionality	Manages HR managers, field managers (by job type and unique email), Records approved jobs (by hash).
Access control and HR manager role-based access via a modifier.	
Job approval mechanism	Requires valid ECDSA signatures from the HR manager and designated field manager for job type.
Data storage and mapping for field managers (job type), job approval (hash), arrays for HR Managers, and job types.	
Events emitted and tracked job approval and HR/field manager additions/removals/updates for auditing.	

4.5.2. Operational Workflow

The workflow follows the following sequential steps:

1. Job Requirement Formulation: HR personnel initiate the process by drafting comprehensive job requirements, including the job title, a detailed description of responsibilities, and the relevant job category.
2. Domain Expert Evaluation: A designated subject-matter expert with pertinent domain expertise meticulously evaluates the technical feasibility and appropriateness of the drafted job posting.
3. Cryptographic Endorsement: Upon satisfactory review, both the responsible HR personnel and the designated domain expert cryptographically sign the finalized job proposal using their private keys.
4. On-Chain Verification and Recording: The Solidity smart contract autonomously verifies the authenticity and validity of the provided digital signatures. Upon successful verification, the contract records the approval of the job posting on the blockchain, ensuring an immutable audit trail. The smart contract is able to be deployed and operate on any Ethereum-based blockchain, not only the Ethereum main net, either public, such as Fantom (https://fantom.foundation/ (accessed on 15 January 2026)), or private/protected, as is the case of Hyperledger Besu (https://besu.hyperledger.org/ (accessed on 15 January 2026)).

4.5.3. Smart Contract Implementation

The complete implementation of the smart contract in Solidity, the programming language for Ethereum smart contracts, along with comprehensive test cases, is publicly available in a dedicated GitHub repository: https://github.com/ValdoMpinga/clean-job (accessed on 15 January 2026).

The key features incorporated within the smart contract implementation include the following:

- Gas-Efficient Signature Recovery: Optimized utilization of the `ecrecover` precompiled contract to minimize the computational cost (gas) associated with signature verification.
- Duplicate Submission Prevention: Implementation of job hashing mechanisms to generate unique identifiers for each job posting, thereby preventing the submission and approval of identical job specifications.

- Event-Driven Architecture: Design of incorporating event emitters that trigger upon significant state changes (e.g., job approval, rejection), facilitating seamless off-chain monitoring and integration with external systems.

Figure 3 presents a visual representation of the data model underpinning the Hardhat job approval smart contract. This model outlines the key entities managed by the contract, including HR managers, field managers, job types, and approved jobs, along with their respective attributes and relationships established between them. The smart contract employs mappings and arrays to maintain and access these data on the Ethereum blockchain, ensuring data integrity and facilitating the job approval process.

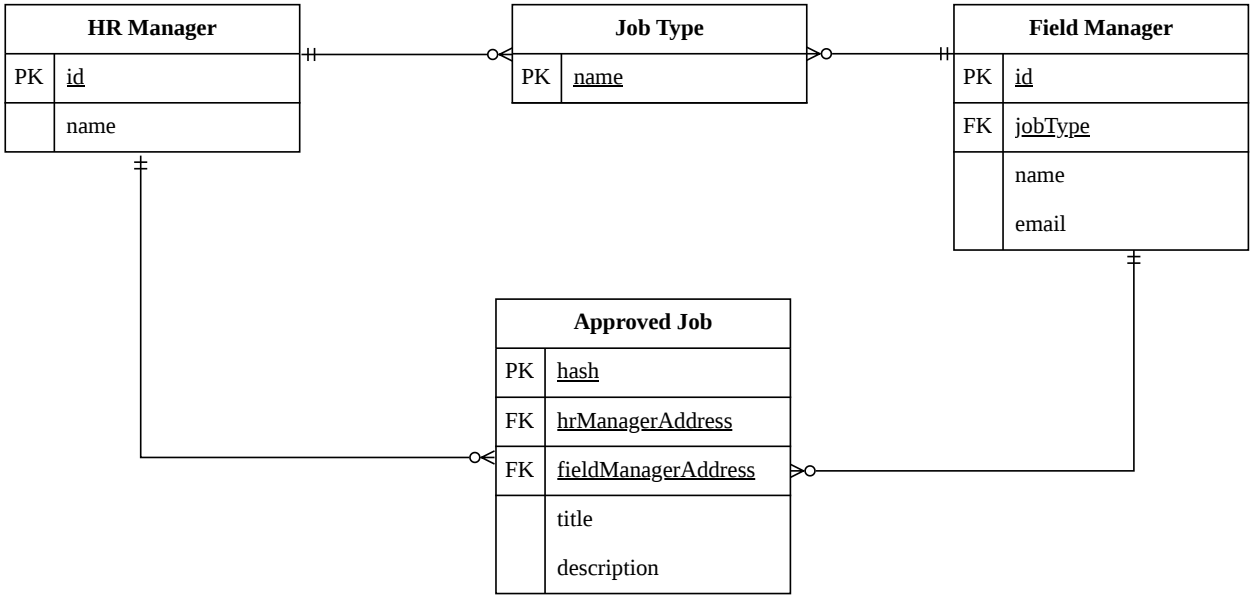


Figure 3. Data model of the Hardhat job approval smart contract.

4.5.4. System Integration and Workflow Phases

The integration and operational lifecycle of the blockchain-based job validation system are structured in two distinct yet interconnected phases:

- The initial smart contract deployment and system configuration.
- The ongoing validation process for individual job postings.

The critical workflow is visually represented using the Business Process Model and Notation (BPMN) diagram in Figure 4. The core concept is to enable each interested entity or company to establish a local copy of the blockchain and operate it within its organizational scope. Consequently, the BPMN diagram has been designed to reflect this decentralized operational model.

4.5.5. Smart Contract Deployment Phase

The initial deployment of the smart contract necessitates the designation of at least one authorized HR manager during the contract's instantiation. On deployment, the HR manager is granted administrative privileges that enable:

- HR Personnel Onboarding: Enrolling additional HR personnel into the system using the `addHRManager()` function of the smart contract.
- Domain Expert Registration: Registering qualified domain experts within the system, associating them with specific job-type taxonomies, using the smart contract function `addFieldManager(jobType)`.

- Job-Type Taxonomy Configuration: Defining and managing the various categories of job types recognized by the system (e.g., “Software Engineering” and “Biomedical Engineering”).

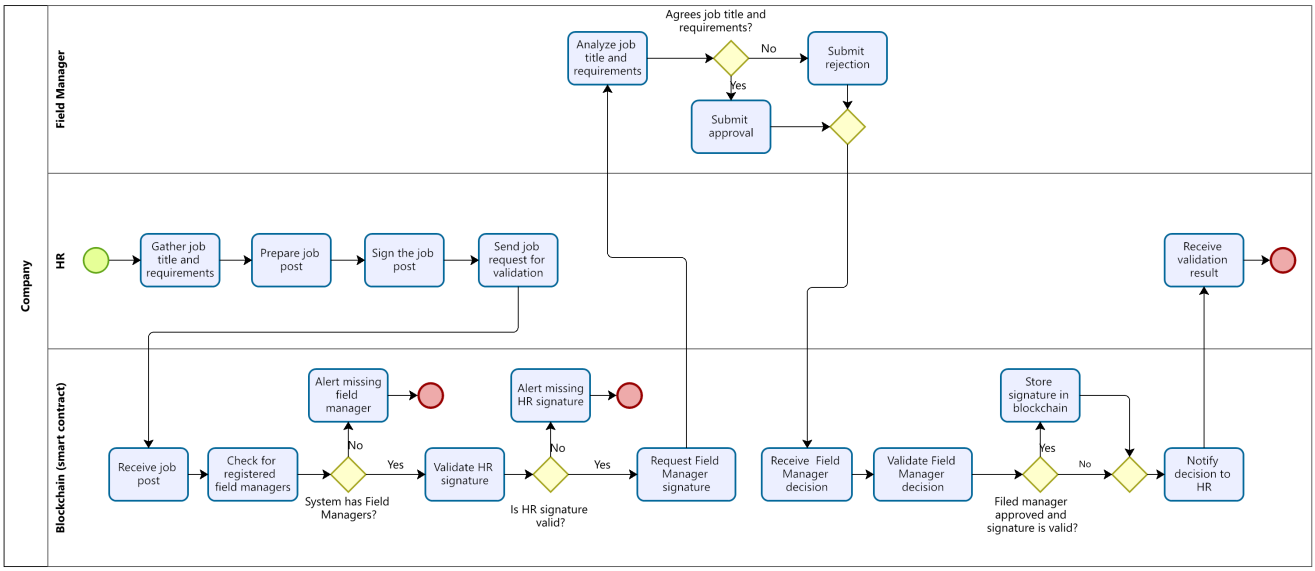


Figure 4. Detailed job posting validation workflow utilizing dual cryptographic signatures.

4.5.6. Job Posting Validation Phase

Figure 4 provides a detailed depiction of the cryptographic validation protocol employed for each job post submitted to the system.

1. Requirement Drafting: HR personnel initiate the process by drafting the complete job specifications, including the title, a comprehensive description of the role, and the designated job type.
2. Expert Assignment and Verification: The system automatically verifies the existence of a registered field expert who is associated with the specified job type. If no qualified expert is currently registered for the given job type, the job posting transaction is automatically reverted, preventing further processing until an appropriate expert has been onboarded.
3. Dual-Signature Validation Flow: Both the responsible HR personnel and the designated field expert independently generate a digital signature for the cryptographic hash of the job posting details. The smart contract then verifies the authenticity of both submitted signatures using the `ecrecover` precompiled contract.
4. Immutable On-Chain Recording: Upon successful verification of both signatures, the approved job posting is securely stored within the `approvedJobs` mapping on the blockchain, creating an immutable record. In instances where the signature verification fails or other validation criteria are not met, the smart contract emits a `JobRejected` event, signaling the rejection of the proposal.

This robust two-phase approach to job validation ensures the following critical properties.

- Enhanced Accountability: All actions performed within the system, including the drafting and approval of job postings, are immutably linked to the cryptographic identities of the participating HR personnel and domain experts.
- Automated Fail-Safety: The smart contract incorporates automated checks and verifications, causing transactions to revert in the event of invalid states or unmet criteria, thereby ensuring the integrity of the validation process.

- Comprehensive Auditability: The inherent transparency and immutability of the blockchain provide a complete and auditable history of all job posting approvals and rejections, fostering trust and accountability within the hiring process.

5. Validation and Discussion

A suite of targeted tests and practical system demonstrations were conducted to rigorously assess the efficacy and robustness of the proposed solution. The next subsections detail the methodologies and outcomes of the evaluations.

5.1. Validation of Job Requirements for Publication

To assess the efficacy of the job vacancy requirement validation service, an expanded data-driven evaluation was conducted using a substantial corpus of real-world job vacancy listings. The publicly accessible "US Jobs on Monster.com" Kaggle dataset, a comprehensive collection encompassing 22,000 job posts originating from the United States, served as the foundational data source for this analysis. From the total dataset of 22,000 job posts, 21,701 were successfully processed and analyzed by the validation service.

The computational infrastructure employed for this evaluation consisted of a domestic desktop system featuring an AMD Ryzen 7 770X processor, 64 GB DDR5 RAM, and an NVIDIA GeForce RTX 4070 graphics processing unit. The validation service was executed within a Windows Subsystem for Linux (WSL) environment running Ollama 3.1:8b. The analysis of the 21,701 job posts was completed in less than 40 h. This timeframe, which is indicative of efficient processing, suggests the potential underutilization of dedicated GPU resources. Monitoring during execution revealed a GPU utilization rate consistently below 10%, indicating a possible bottleneck or configuration issue within the WSL environment that limits the full exploitation of the parallel processing capabilities of RTX 4070. Despite this observed limitation, the total processing time remains good for the scale of the analyzed dataset. The results of these samples highlight a significant dehumanization issue in today's job listings, as evidenced by the following findings (see Figure 5):

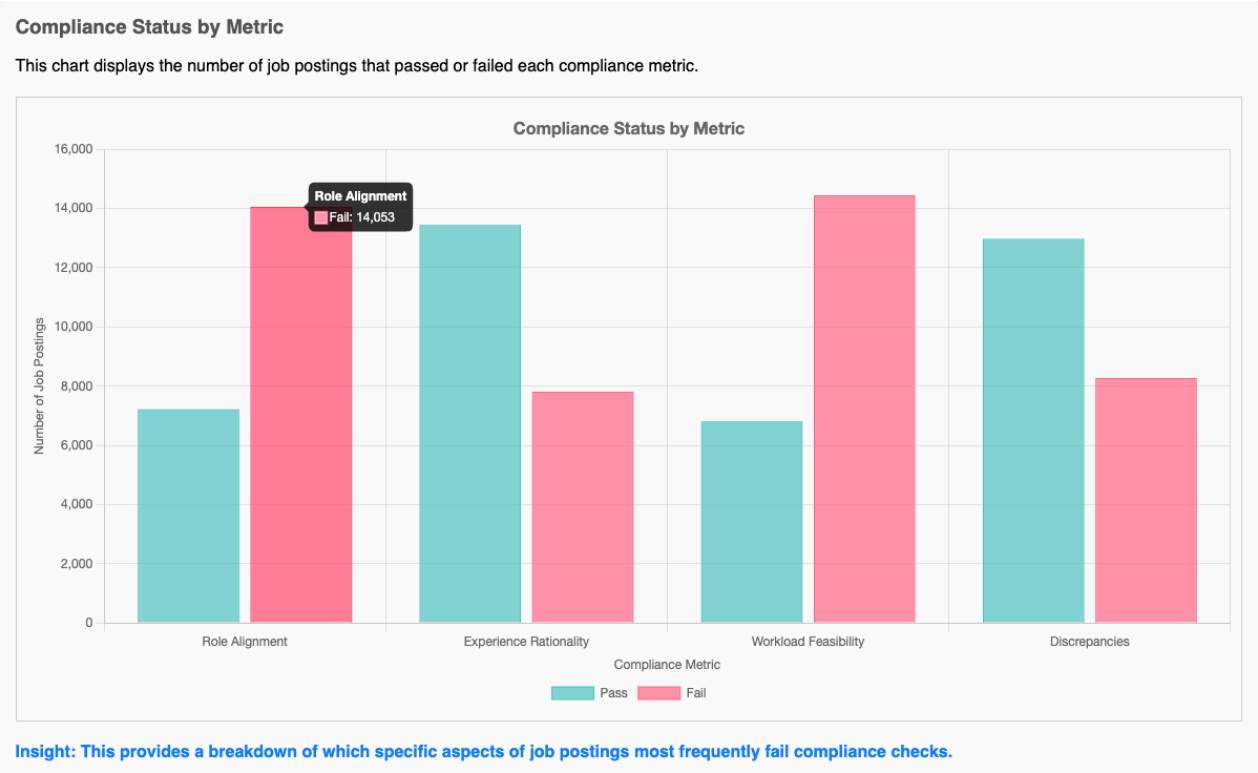


Figure 5. Evaluation stats for the vacancy validator service.

- In total, 64.76%, with a 95% confidence interval (CI) between 64.12 and 65.39%, out of the analyzed 21,701 job postings failed our role alignment check (14,053 out of 21,701).
- A total of 35.99% of the jobs, with 95% CI between 35.35 and 36.63%, had unreasonable experience requirements (7816 out of 21,701 failed Experience Rationality).
- In total, 66.56% of the workloads, with 95% CI between 65.93 and 67.19%, were unrealistic for the role of one person (14,445 out of 21,701 failed Workload Feasibility).
- A total of 38.11% of the jobs, with 95% CI between 37.46 and 38.76%, contained discrepancies (8270 out of 21,701).

These values should be interpreted as module-generated classifications on the analyzed dataset and not as externally verified estimates of the prevalence of problematic job postings.

5.2. Revealing Bias Triggers

The service endpoint was determined as expected. Typically, ATSs extract user data from a CV and generate outputs in formats such as JSON, XML, or plain text. To test this functionality, mock data were generated in the format shown in the following example:

```
{
  "name": "John Doe",
  "age": 29,
  "gender": "Male",
  "sexual_orientation": "Heterosexual",
  "race": "Caucasian",
  "religion": "Christian",
  "disability_status": "None",
  "marital_status": "Single",
  "children": 0,
  "languages_spoken": ["English", "Spanish"],
  "photo": "http://example.com/photo.jpg",
  "address": "123 Main St, Cityville, USA",
  "email": "johndoe@example.com",
  "phone": "123-456-7890"
}
```

The expected output for these data is:

```
{
  "languages_spoken": ["English", "Spanish"],
  "address": "123 Main St, Cityville, USA",
  "email": "johndoe@example.com",
  "phone": "123-456-7890",
}
```

The functional evaluation of the module was conducted on a set of structured mock applicant records generated according to the same schema, with the purpose of verifying whether bias-conducive fields could be identified and removed while preserving data integrity. This evaluation should be interpreted as a controlled validation of the module's behavior under the defined input format, rather than as a full validation on real-world CV or ATS datasets.

The module's endpoint successfully produces this output for JSON, XML, and plain-text data.

Testing the module on real or realistically annotated CV/ATS corpora is an important direction for future work

5.3. Evaluation of Different Large Language Models by Number of Parameters Using Job Requirement Analysis

Following the initial implementation utilizing the LLaMA model, further investigation was conducted to evaluate the performance of a range of other open-source LLMs with varying parameter counts for the task of job requirement analysis.

The purpose of this evaluation was not to establish a universally optimal LLM across all recruitment-related subtasks but to assess the operational suitability of candidate models to produce structured, pattern-like output, a crucial aspect for automating the analysis of multiple job requirements, and for integration into the proposed Humanization Services artifacts.

The selected indicators were parsing success rate, compliance score, and processing time. These were chosen because they directly reflect the practical requirements of the implemented pipeline, namely, the ability to generate valid structured outputs, conform to the expected response schema, and operate with acceptable responsiveness. Therefore, the evaluation should be interpreted as an artifact-oriented assessment of feasibility and applicability under the tested conditions.

5.3.1. Large Language Model Parameters and Significance (In the Context of Broader LLM Understanding)

As previously discussed, the number of parameters is a key characteristic of LLMs, influencing their ability to learn and generate text [35]. While larger models generally exhibit enhanced capacity for complex language understanding, they also demand greater computational resources. This evaluation explores this relationship across a spectrum of open-source models tailored for different resource constraints.

5.3.2. Selected Large Language Models

To evaluate the performance of LLMs in job requirement analysis, we selected a range of open-source models categorized into three size groups to represent different computational resource constraints (see Table 3).

Table 3. Evaluation of large language models for job requirement analysis.

Model Category	Name	Time (s)	Role Align	Exp. Ratio	Work-Load	Dis-Crep.	Compl.	Output Issues
Ultra-light	qwen2:0.5b	2.37	N/A	N/A	N/A	N/A	N/A	Could not parse.…
Ultra-light	qwen2:0.5b	1.19	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Ultra-light	qwen2:0.5b	0.27	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Ultra-light	llama3.2:1b	4.15	Pass	Fail	Pass	Pass	50	…
Ultra-light	llama3.2:1b	1.86	Pass	Fail	Pass	N/A	Missing 'summary' key…	
Ultra-light	llama3.2:1b	1.95	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Ultra-light	gemma3:1b	5.31	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Ultra-light	gemma3:1b	4.35	Pass	Fail	Pass	Pass	48/100	…
Ultra-light	gemma3:1b	3.38	Pass	Fail	Pass	Pass	65/100	…
Light	phi3:mini	14.01	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Light	phi3:mini	4.59	N/A	N/A	N/A	N/A	N/A	Parsing Error ….
Light	phi3:mini	3.84	Pass	Pass	Fail	Fail	60	…
Light	llama3.2:3b	10.79	Fail	Pass	Fail	Pass	60	…

Table 3. *Cont.*

Model Category	Name	Time (s)	Role Align	Exp. Ratio	Work-Load	Dis-Crep.	Compl.	Output Issues
Light	llama3.2:3b	3.47	Fail	Pass	Fail	Pass	60	…
Light	llama3.2:3b	2.68	Fail	Pass	Fail	Fail	0	…
Light	gemma3:4b	14.65	Pass	Fail	Fail	Pass	75	…
Light	gemma3:4b	7.78	Pass	Fail	Fail	Pass	40	…
Light	gemma3:4b	7.11	Pass	Fail	Fail	Pass	65	…
Medium-light	mistral:7b	13.31	Pass	Pass	100	…		
Medium-light	mistral:7b	5.97	Pass	Fail	Pass	Fail	40	…
Medium-light	mistral:7b	6.16	Pass	Pass	Fail	Pass	80	…
Medium-light	qwen2:7b	14.14	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Medium-light	qwen2:7b	1.32	N/A	N/A	N/A	N/A	N/A	Parsing Error: ….
Medium-light	qwen2:7b	5.98	Pass	Fail	Pass	Fail	50	…
Medium-light	llama3.1:8b	17.86	Fail	Pass	Fail	Pass	70	…
Medium-light	llama3.1:8b	5.50	Fail	Pass	Pass	Pass	67	…
Medium-light	llama3.1:8b	6.25	Pass	Fail	Pass	Fail	40	…

5.3.3. Analysis of LLMs' Comparison Results

Output Parsing Errors

A significant observation from Table 3 is the prevalence of parsing errors, particularly with the smaller models. Models such as Qwen2-0.5B consistently failed to produce valid JSON outputs across all test runs, with 100% failure rate, rendering their evaluations unusable. Similarly, Phi-3-Mini frequently struggled with output formatting, showing a 67% parsing failure rate despite being in the light category. This indicates that these models, although efficient in terms of computational cost, may struggle with the precision required for structured output generation. Parsing errors are likely due to the model's tendency to include extraneous text or deviate from the specified JSON format, suggesting insufficient training for structured data tasks.

Processing Time

As expected, there is a general trend of increased processing time with larger models. The ultra-light models (e.g., Qwen2-0.5B with times ranging from 0.27 s to 2.37 s, Llama-3.2-1B from 1.86 s to 4.15 s) exhibited the fastest processing times, often below 5 s per job posting. The light models showed increased variability, with Phi-3-Mini ranging from 3.84 s to 14.01 s, while the medium-light models (e.g., Mistral-7B from 5.97 s to 13.31 s, Llama-3.1-8B from 5.50 s to 17.86 s) generally took longer, with some exceeding 10 s. Notably, the processing time variance within individual models suggests inconsistent optimization or varying complexity handling capabilities.

Evaluation Accuracy and Consistency

Among the models that produced parsable output, distinct patterns emerged across the evaluation criteria. **Role alignment** showed mixed results, with Mistral-7B demonstrating consistent performance, while Llama models (both 3.2:3B and 3.1:8B) frequently failed this criterion. **Experience ratio** evaluation proved to be the most challenging aspect across all model categories, with failures observed in 78% of successful parsing attempts, indicating a fundamental difficulty in assessing the reasonableness of experience requirements. This result may reflect at least two factors. First, the relationship between job seniority and years of experience is often context-dependent and may not be explicitly represented in

the prompt or source text. Second, the prompt structure may not sufficiently constrain the model to apply a consistent reasoning pattern when evaluating experience requirements. This suggests that future improvements should focus on role-specific validation heuristics, clearer prompt decomposition, and more explicit rules linking seniority levels to expected experience ranges. **Workload Feasibility** and the **Discrepancy Check** showed more balanced performance, particularly in larger models. The compliance scores varied significantly, ranging from 0 to 100, with Mistral-7B achieving the highest score of 100 in one instance and maintaining scores between 40 and 80 in other runs.

Model Capacity and Complexity

The results suggest that larger models with higher parameter counts are better equipped to handle the complexity of the job requirement analysis task. However, parameter count alone does not guarantee superior performance. Qwen2-7B, despite its larger size, showed significant parsing failures (67% failure rate), while smaller models like Gemma-3-1B achieved better parsing success rates (67% success). This indicates that model architecture, training methodology, and optimization play crucial roles beyond raw parameter scaling. The task requires nuanced understanding of language, the ability to identify subtle biases, and the capacity to reason about the implications of different requirements—capabilities that appear to emerge more reliably in well-optimized models around the 7B parameter range.

Performance Hierarchy and Practical Implications

Based on the comprehensive evaluation, a clear performance hierarchy emerged. Mistral-7B demonstrated the most promising overall performance, combining reliable parsing (100% success rate), balanced analytical capabilities across all criteria, and reasonable processing times. Gemma-3-4B showed consistent parsing success but analytical limitations, particularly in experience ratio evaluation. The ultra-light models, while computationally efficient, proved insufficient for reliable structured output generation in this domain. These findings suggest that for practical deployment, a minimum model size of approximately 3-4B parameters is required for parsing reliability, with 7B+ parameters needed for robust analytical performance.

Implications for Humanization Potential

Despite the challenges with output formatting and consistency in some models, the overall feasibility of using LLMs to automate job requirement analysis remains promising. Models such as Mistral-7B have demonstrated the potential to provide valuable insights into the fairness and rationality of job postings, achieving compliance scores that indicate meaningful analytical capability. The consistent weakness in experience ratio evaluation across all models highlights an area requiring further development, potentially through domain-specific fine-tuning or enhanced reasoning frameworks. This technology can assist human reviewers in identifying potential issues, ultimately contributing to a more human-centered recruitment process, though current limitations suggest the need for human oversight, particularly in experience-related assessments.

5.4. Validating the Blockchain-Based Validation System

The core validation logic within the blockchain-based validation system's smart contract was subjected to a rigorous verification process using a test suite developed with Hardhat for the testing environment and Chai for assertions. This comprehensive suite specifically targets three critical aspects fundamental to the system's security and operational integrity.

5.4.1. Access Control Testing

- HR Manager Exclusivity: The tests successfully confirmed that the system strictly enforces HR manager exclusivity for administrative functions, effectively rejecting all unauthorized attempts to modify critical system parameters or roles by non-HR actors.
- Field Expert Assignment Validation: The test suite validated the system's ability to enforce the assignment of field experts based on specific job categories, ensuring that only designated experts are authorized to provide validation for relevant job postings.
- Secure Role Revocation: Functionality for the revocation of assigned roles was tested and confirmed to operate correctly without causing any corruption or inconsistencies in the system's internal state.

5.4.2. Robust Signature Verification Testing

- Invalid Signature Detection: In the executed automated test cases, the system correctly rejected all job approval attempts containing invalid HR or field manager digital signatures, showing that only cryptographically authorized personnel could endorse job postings (see Figure 6).
- Unauthorized Approval Prevention: Within the tested scenarios, the ECDSA (Elliptic Curve Digital Signature Algorithm) validation mechanism prevented all unauthorized job approval attempts by entities that lacked valid digital signatures.
- Duplicate Submission Blocking: In the automated tests, the implemented job-hashing mechanism successfully blocked duplicate submission and approval attempts, thereby preserving the integrity and uniqueness of validated vacancies.

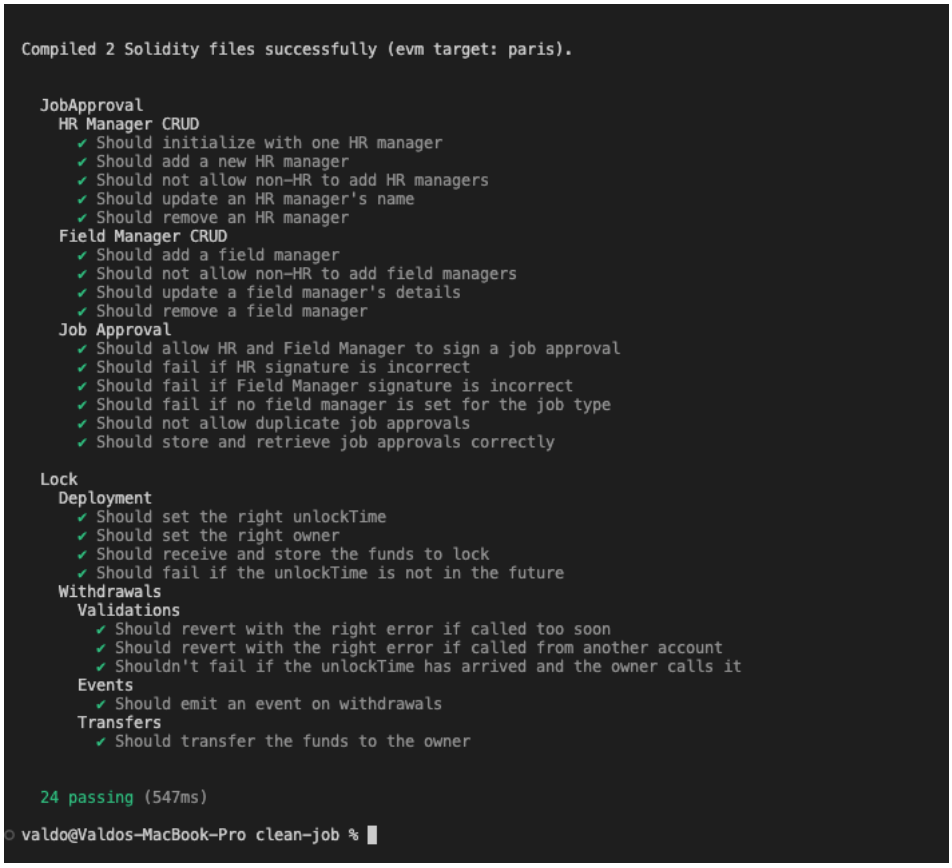


Figure 6. Automated test results showing 24 passing tests for the blockchain-based validation module.

5.4.3. End-to-End Workflow Integrity Testing

The designed blockchain-based job post validation tool ensures the involvement of human actors (HR manager and field manager) before the job post is deployed. Keeping the human in the loop, demands human accountability and increases job requirements alignment with job function and recruitment fairness.

- Mandatory Expert Assignment Enforcement: The tests confirmed that the system enforces the mandatory assignment of a qualified field expert for a given job category before the validation process can be initiated, ensuring that all job postings receive appropriate domain-specific review.
- Consistent State Management: The test suite verified that the system maintains a consistent and accurate internal state across all Create, Read, Update, and Delete (CRUD) operations related to HR managers, field experts, and job postings.
- Approval History Immutability: The tests confirmed that the approval history of job postings, once recorded on the blockchain, remains immutable and cannot be retroactively altered or tampered with after transaction finalization.

The collective results of this comprehensive test suite provide strong empirical evidence that the blockchain validation system correctly and securely implements the following.

- Robust role-based access control mechanisms, ensuring that only authorized entities can perform specific actions.
- Rigorous cryptographic requirements validation through ECDSA signatures, guaranteeing the authenticity and integrity of approvals.
- Full adherence to all functional requirements as originally outlined in the smart contract's design specifications.

In Appendix A, a system demonstration via the public web interface can be seen.

5.5. Limitations and Future Work

At this point, several study limitations must be acknowledged.

In the LLM evaluation, although the reported indicators are adequate for assessing the operational suitability of the models for integration into the proposed Humanization Service, they do not capture all dimensions relevant to a comprehensive benchmark-oriented comparison. In particular, future work should extend the evaluation with finer-grained task-level accuracy metrics, measures of sensitive-information omission and valid-information retention, repeated-run consistency analysis, structured-output syntax error rates, and infrastructure-level resource consumption indicators. Such extensions would provide a more detailed basis for model selection in large-scale or production-oriented settings.

The empirical comparison reported in this study reflects the models available during the study period and should therefore be interpreted as an evaluation of model suitability within the implemented artifact context, rather than as an exhaustive benchmark of all subsequently released LLMs.

Despite a demonstration of the designed and implemented platform being presented, neither a user validation study nor user acceptance tests have been developed. Future work should focus on validation and user acceptance testing with HR professionals to gather their feedback on how and to what extent the developed tool improves fairness and the alignment of job specifications with actual job requirements.

Also, having found no reliable reports or published humanization approaches with comparable metrics, namely, percentage of job descriptions failing role alignment, experience rationality and workload feasibility and other discrepancies, it was not possible to make a benchmark comparison. Besides the lack of benchmarking, the study only includes

limited error analysis and confidence intervals. Future work also needs to improve the statistical significance of the study.

6. Conclusions

The dehumanization of job roles and the risk of violating fundamental human rights in recruitment remain significant ethical, organizational, and technical concerns. As AI-based tools become increasingly embedded in ATSs and recruitment workflows, there is a growing need for approaches that improve fairness, transparency, and accountability without sacrificing the operational benefits of automation. Recruitment systems may reproduce or amplify existing biases through problematic job requirements, opaque screening criteria, and insufficient human oversight.

This study addressed these challenges through the design, implementation, and evaluation of a set of Humanization Services intended to intervene at critical points of the recruitment pipeline. Following a Design Science Research approach, three main artifacts were developed: a Job Requirements Validation Module, a Bias Trigger Removal Module, and a blockchain-supported dual-authorization mechanism for vacancy approval. Together, these components were designed to improve the quality of job postings, reduce bias-conducive information in applicant data, and strengthen human accountability in the publication of job vacancies.

The results provide initial empirical support for the operational feasibility of the proposed approach under the tested conditions. The Job Requirements Validation Module identified problematic vacancy characteristics, including misalignments, unrealistic expectations, unreasonable demands, and internal inconsistencies. The Bias Trigger Removal Module showed that applicant data can be processed in a way that reduces exposure to sensitive or bias-conducive attributes while preserving usability for subsequent stages of the workflow. The blockchain-based approval mechanism showed that dual human authorization can be technically enforced through cryptographic validation, thereby supporting traceability, non-repudiation, and accountability in vacancy publication.

These findings should not be interpreted as evidence that the proposed approach fully solves bias in recruitment or demonstrates broad organizational impact beyond the evaluated artifacts. Rather, the study contributes a practical and theoretically grounded step toward more human-centered ATS-supported recruitment by showing that fairness-oriented interventions can be embedded in recruitment workflows at the level of vacancy preparation, applicant-data handling, and human approval processes.

The study further highlights that humanization in recruitment should be understood as a socio-technical and systemic challenge rather than a purely algorithmic one. Fairness in ATS-supported recruitment is not a property of a single model or decision point but an emergent outcome of the interaction among job postings, applicant data, automated screening components, human reviewers, and governance mechanisms. In this sense, the proposed Humanization Services contribute not only as technical artifacts but as socio-technical interventions intended to improve the alignment between automation and organizational values such as fairness, transparency, accountability, and human dignity.

From a systems-theoretic and socio-technical perspective, the main contribution of this study is to show that fairness in ATS-supported recruitment cannot be treated as a property of a single model or decision point. Rather, recruitment should be understood as an open organizational system in which job postings, applicant data, automated screening components, human reviewers, and governance mechanisms interact to shape outcomes. In this sense, the proposed Humanization Services contribute not only as technical artifacts but as socio-technical interventions that seek to improve the alignment between automation and organizational values such as fairness, transparency, accountability, and human dignity.

By addressing job posting integrity, bias-conducive applicant data, and enforceable human oversight within the same workflow, the study supports the view that bias in recruitment is best approached as an emergent systemic phenomenon requiring coordinated technical and organizational responses.

Several limitations remain and define important directions for future work. Additional empirical validation with HR professionals and domain experts is needed to assess how the proposed approach performs in real organizational settings and how it affects fairness, usability, and trust in practice. Further research should also develop benchmark datasets and broader evaluation criteria for comparing humanization-oriented recruitment approaches. In addition, future work may explore candidate-centered extensions, such as transparent feedback mechanisms, while carefully evaluating their ethical implications and practical effectiveness. Since the system operates on potentially sensitive personal data, a deeper analysis of compliance with GDPR and other data-protection frameworks is likewise necessary.

Author Contributions: Conceptualization, V.V.M.; methodology, V.V.M. and A.M.R.d.C.; software, V.V.M.; validation, V.V.M. and A.M.R.d.C.; formal analysis, V.V.M. and A.M.R.d.C.; investigation, V.V.M.; writing—original draft preparation, V.V.M. and A.M.R.d.C.; writing—review and editing, A.M.R.d.C.; visualization, V.V.M.; supervision, A.M.R.d.C.; project administration, A.M.R.d.C.; funding acquisition, A.M.R.d.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by The Applied Digital Transformation Laboratory (ADiT-LAB), through the Portuguese Foundation for Science and Technology (FCT - Fundação para a Ciência e a Tecnologia) within project: UIDP/06121/2025 and DOI identifier https://doi.org/10.54499/UID/06121/2025.

Institutional Review Board Statement: This research study uses anonymous human data publicly available with a CC BY-SA 4.0 license (https://creativecommons.org/licenses/by-sa/4.0/) (accessed on 15 January 2026) on https://www.kaggle.com/datasets/PromptCloudHQ/us-jobs-on-monstercom (accessed on 15 January 2026).

Data Availability Statement: The source code of the artifacts presented in the study are openly available in https://github.com/ValdoMpinga/clean-job-UI (accessed on 15 January 2026), https://github.com/ValdoMpinga/recruitment-humanization-service (accessed on 15 January 2026) and https://github.com/ValdoMpinga/clean-job (accessed on 15 January 2026). The data used in this study are publicly available at https://www.kaggle.com/datasets/PromptCloudHQ/us-jobs-on-monstercom (accessed on 15 January 2026).

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT, version 5.4, for the purposes of text improvement, summarization and consolidation. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
API	Application Programming Interface
ATS	Application Tracking System
BCT	Blockchain Technology
BPMN	Business Process Model and Notation
CI	Confidence Interval
CV	Curriculum Vitae
DSR	Design Science Research

ECDSA	Elliptic Curve Digital Signature Algorithm
HR	Human Resources
HRM	Human Resources Management
JSON	JavaScript Object Notation
LLM	Large Language Model
ML	Machine Learning
NLP	Natural Language Processing
UI	User Interface

Appendix A. System Demonstration via Public Web Interface

This interactive platform implements three core functionalities, providing tangible demonstrations of the system's capabilities.

Appendix A.1. Interactive Job Requirement Analysis

- **Accessibility**: The job requirement analysis tool is publicly accessible via the following URL: https://joblimpo.valdompinga.com/requirements (accessed on 15 January 2026).
- **Real-Time Evaluation**: This interface enables users to perform real-time evaluations of job postings against a predefined set of human-centered criteria, providing immediate feedback on potential issues.
- **Natural Language Processing**: The tool leverages the underlying validation framework to process natural language input from job postings, identifying and highlighting areas of concern based on the defined metrics (see Figure A1).

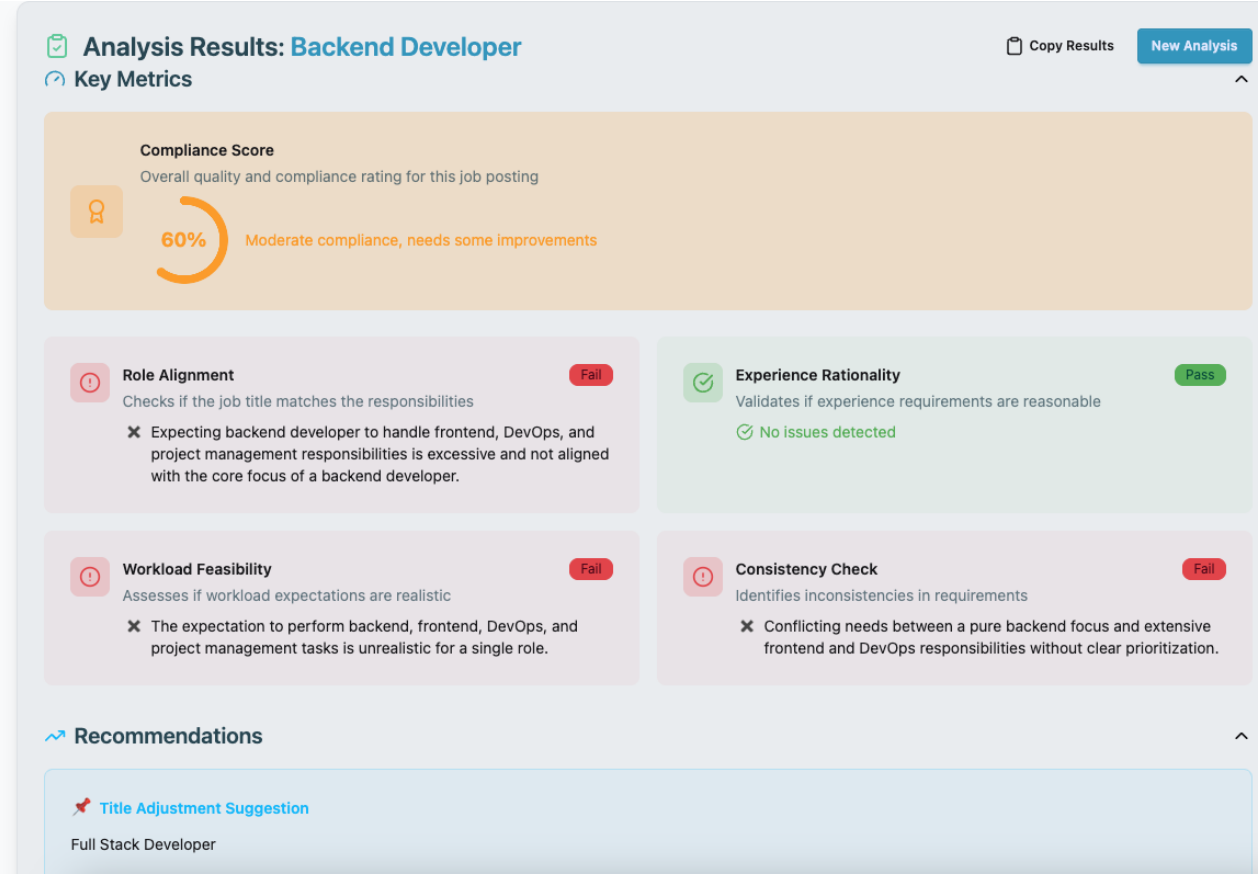


Figure A1. Interactive web interface for job requirement evaluation, displaying dehumanization detection metrics derived from natural language analysis.

Appendix A.2. Bias-Free Candidate Presentation Interface

- **Accessibility**: The candidate anonymization service demonstration is available at https://joblimpo.valdompinga.com/candidate (accessed on 15 January 2026).
- **Demonstration Across Multiple Data Formats**: This section showcases the practical implementation of the proposed candidate anonymization solution, illustrating its effectiveness in processing candidate data presented in JSON, XML, and plain text formats.
- **Bias Indicator Removal**: The following figures demonstrate how the system effectively identifies and removes sensitive demographic indicators from candidate profiles provided in each of these formats while preserving essential professional qualifications and experience details, thereby mitigating potential unconscious biases.

Appendix A.2.1. JSON Format

Figure A2 illustrates the anonymization process for candidate data provided in the JSON format.

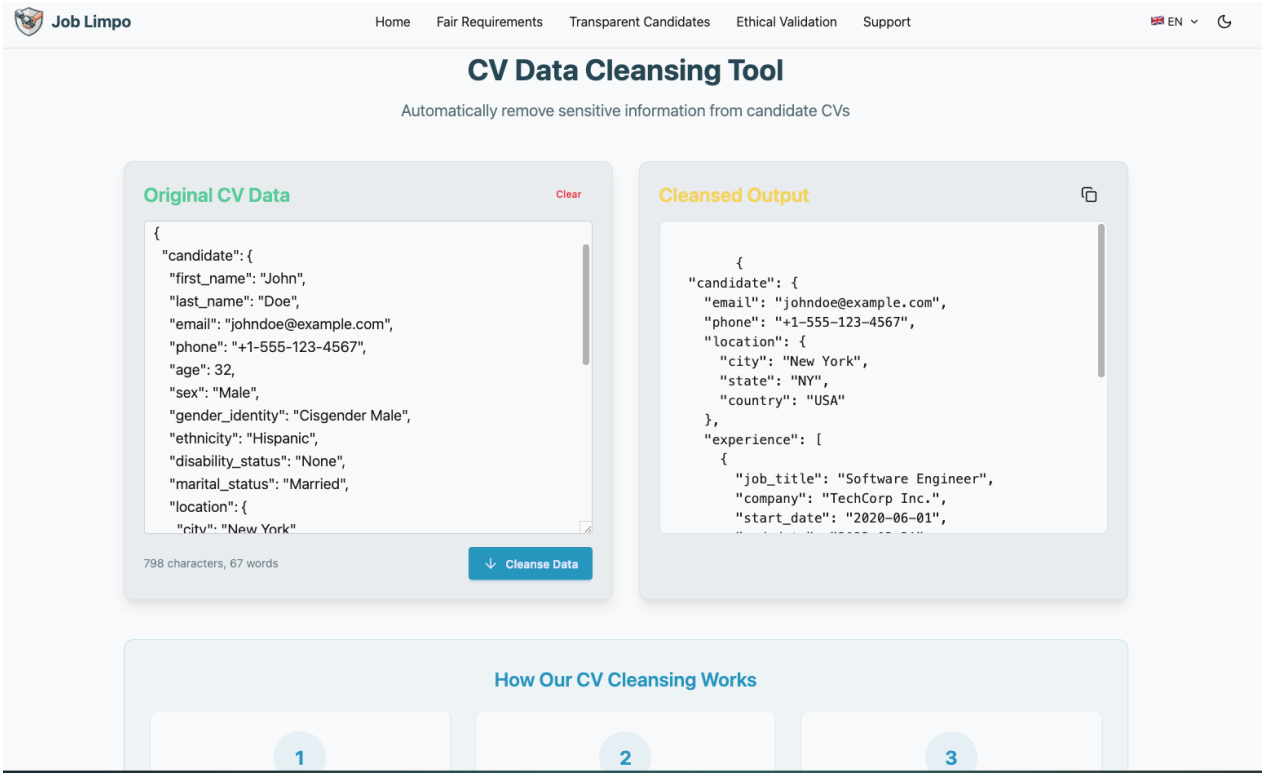


Figure A2. Web interface demonstrating candidate profile processing with bias-inducing information removal for JSON format.

Appendix A.2.2. XML Format

Figure A3 illustrates the anonymization process for candidate data provided in the XML format.

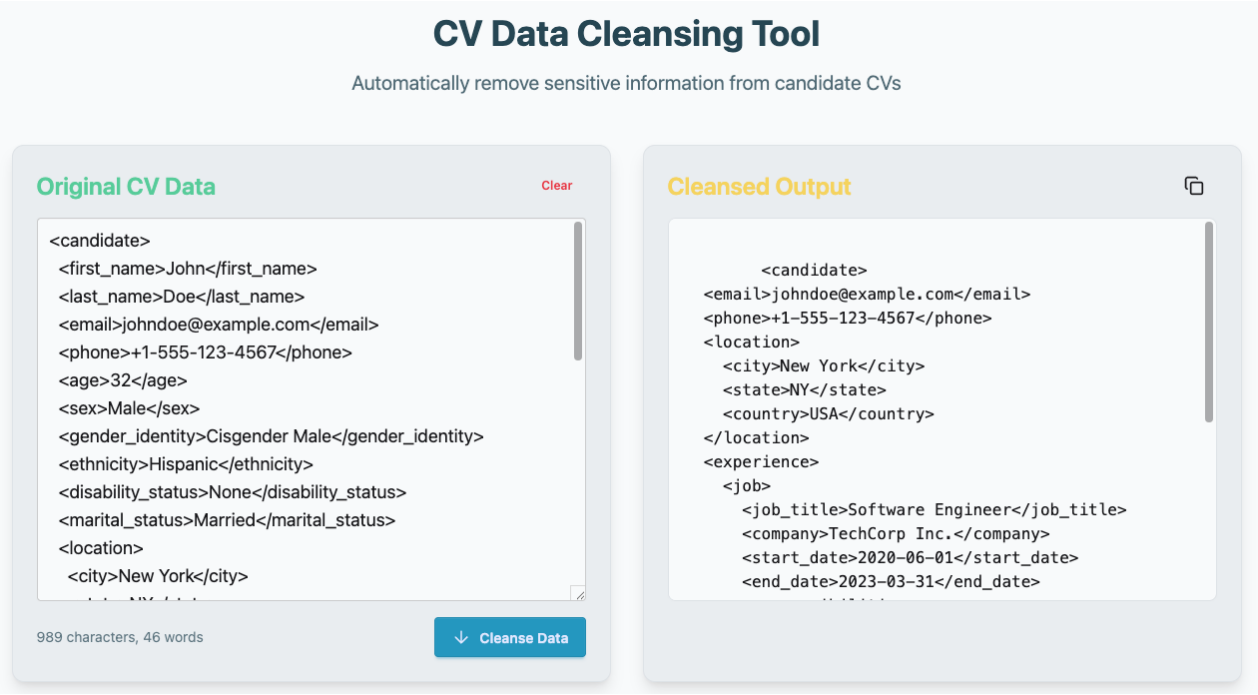


Figure A3. Web interface demonstrating candidate profile processing with bias-inducing information removal for XML format.

Appendix A.2.3. Plain Text Format

Figure A4 illustrates the anonymization process for the candidate data provided in plain-text format.

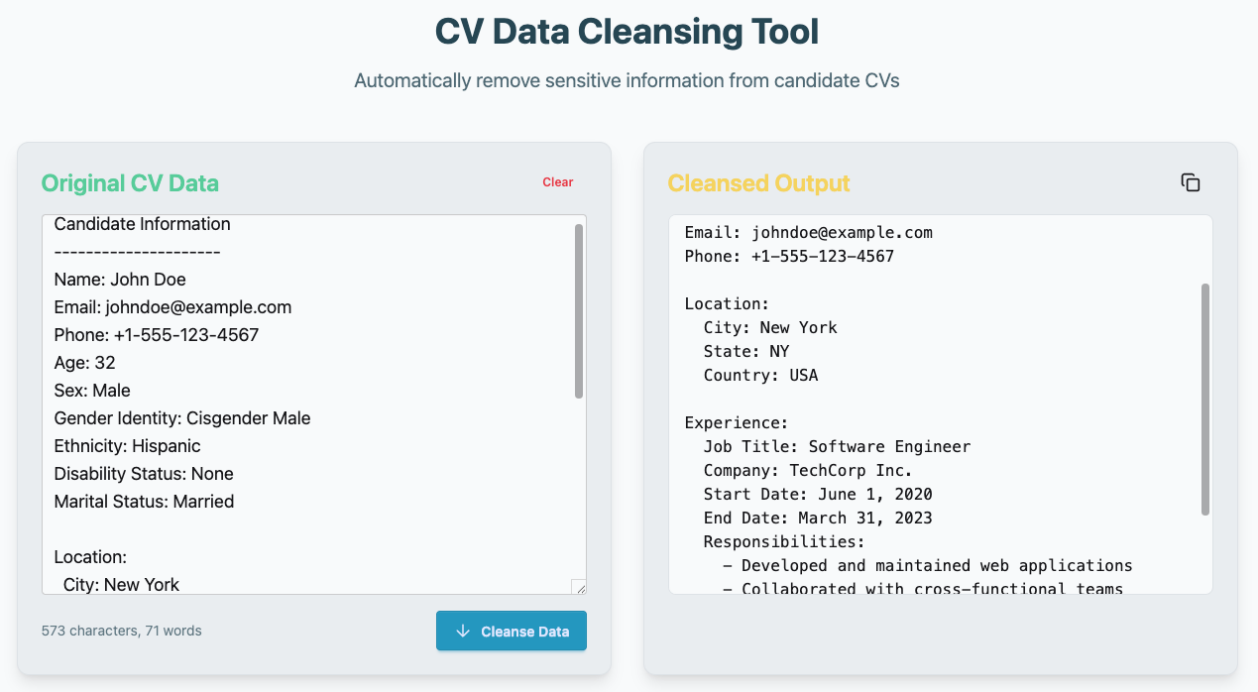


Figure A4. Web interface demonstrating candidate profile processing with bias-inducing information removal for plain text format.

Appendix A.3. Blockchain-Based Service Implementation Showcase

- **Platform Access**: The interface providing access to the blockchain implementation details is hosted at https://joblimpo.valdompinga.com/validation (accessed on 15 January 2026).
- **Open-Source Repository Link**: The project's public GitHub repository is https://github.com/ValdoMpinga/clean-job/ (accessed on 15 January 2026). It contains the following critical components (see Figure A5):
  - The complete Solidity source code for the smart contracts implementing the blockchain validation logic.
  - The comprehensive Hardhat test suite used for rigorous contract verification.
  - Deployment scripts facilitating the deployment and initialization of the smart contracts on the Ethereum network.

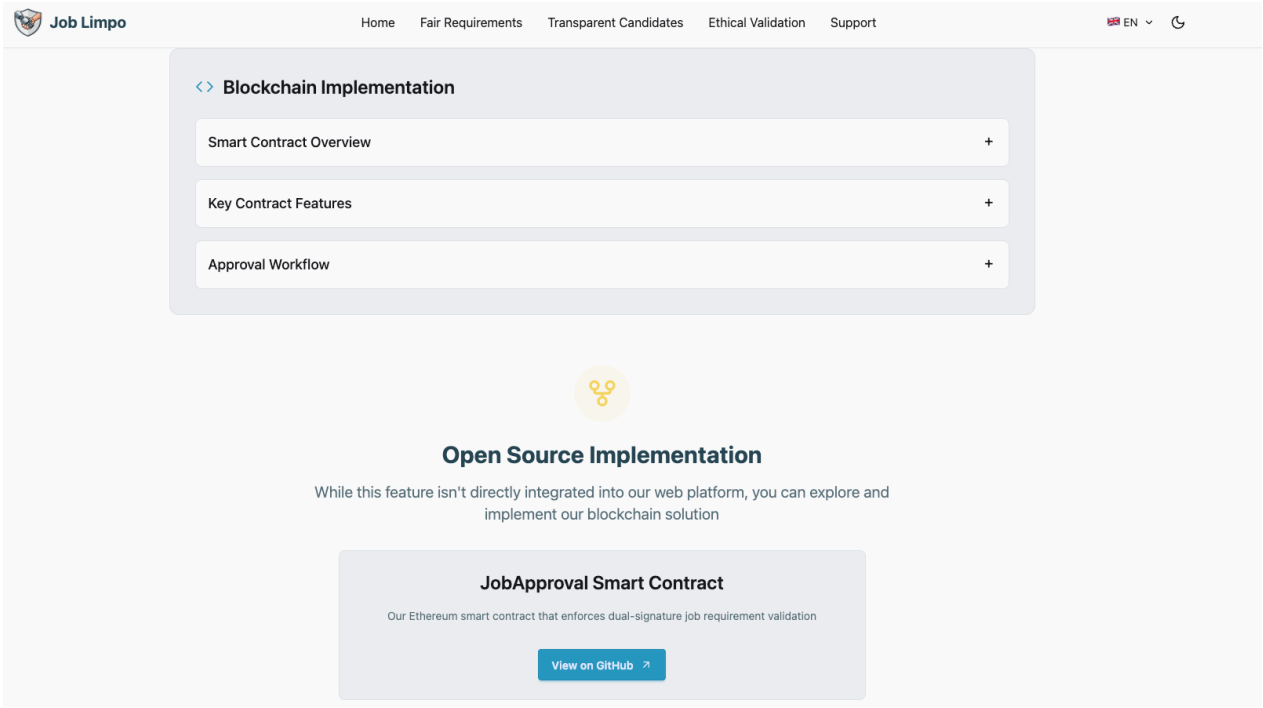


Figure A5. Web interface providing a link to the open-source blockchain solution repository.

References

1. Recruitify.ai. The Evolution of Applicant Tracking Systems (ATS): From Manual Processes to AI-Powered Recruitment. 2023. Available online: https://www.recruitify.ai/en/blog/evolution-of-ats-systems-from-manual-processes-to-ai-driven-recruitment/ (accessed on 1 May 2025).
2. Yam, J.; Skorburg, J.A. From human resources to human rights: Impact assessments for hiring algorithms. *Ethics Inf. Technol.* **2021**, *23*, 611–623. [CrossRef]
3. Human Rights Council. *Promotion and Protection of All Human Rights, Civil, Political, Economic, Social and Cultural Rights, Including the Right to Development, 2020*; Agenda Item 3, Forty-Third Session, 24 February–13 March and 15–23 June 2020; Human Rights Council: Geneva, Switzerland, 2020.
4. Human Rights Council. *Promotion and Protection of All Human Rights, Civil, Political, Economic, Social and Cultural Rights, Including the Right to Development, 2024*; Agenda item 3, Fifty-fifth Session, 26 February–5 April 2024; Human Rights Council: Geneva, Switzerland, 2024.
5. Amnesty International. Surveillance Giants: How the Business Model of Google and Facebook Threatens Human Rights, 2019. Available online: https://www.amnesty.org/en/documents/pol30/1404/2019/en/ (accessed on 23 July 2025).
6. Fabris, A.; Baranowska, N.; Dennis, M.J.; Graus, D.; Hacker, P.; Saldivar, J.; Zuiderveen Borgesius, F.; Biega, A.J. Fairness and Bias in Algorithmic Hiring: A Multidisciplinary Survey. *ACM Trans. Intell. Syst. Technol.* **2025**, *16*, 1–54. [CrossRef]
7. Rigotti, C.; Fosch-Villaronga, E. Fairness, AI & recruitment. *Comput. Law Secur. Rev.* **2024**, *53*, 105966. [CrossRef]

8. Kumar, D.; Grosz, T.; Rekabsaz, N.; Greif, E.; Schedl, M. Fairness of recommender systems in the recruitment domain: An analysis from technical and legal perspectives. *Front. Big Data* **2023**, *6*, 1245198. [CrossRef] [PubMed]
9. Zhang, S.; Kuhn, P.J. *Measuring Bias in Job Recommender Systems: Auditing the Algorithms*; Working Paper 32889; National Bureau of Economic Research: Cambridge, MA, USA, 2024.
10. Del Carmen Fernández Martínez, M.; Fernández, A. AI in recruiting. Multi-agent systems architecture for ethical and legal auditing. *IJCAI Int. Jt. Conf. Artif. Intell.* **2019**, *2019*, 6428–6429. [CrossRef]
11. Aizenberg, E.; Dennis, M.J.; van den Hoven, J. Examining the assumptions of AI hiring assessments and their impact on job seekers' autonomy over self-representation. *AI Soc.* **2023**, *40*, 919–927. [CrossRef] [PubMed]
12. Lavanchy, M.; Reichert, P.; Narayanan, J.; Savani, K. Applicants' Fairness Perceptions of Algorithm-Driven Hiring Procedures. *J. Bus. Ethics* **2023**, *188*, 125–150. [CrossRef]
13. Paramita, D.; Okwir, S.; Nuur, C. Artificial intelligence in talent acquisition: Exploring organisational and operational dimensions. *Int. J. Organ. Anal.* **2024**, *32*, 108–131. [CrossRef]
14. Cai, F.; Zhang, J.; Zhang, L. The Impact of Artificial Intelligence Replacing Humans in Making Human Resource Management Decisions on Fairness: A Case of Resume Screening. *Sustainability* **2024**, *16*, 3840. [CrossRef]
15. Ling, B.; Dong, B.; Cai, F. Applicants' Fairness Perception of Human and AI Collaboration in Resume Screening. *Int. J. Hum.-Comput. Interact.* **2025**, *41*, 10787–10798. [CrossRef]
16. Wilson, K.; Caliskan, A. Gender, Race, and Intersectional Bias in Resume Screening via Language Model Retrieval. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*; The AAAI Press: Washington, DC, USA, 2024; Volume 7, pp. 1578–1590. [CrossRef]
17. Seshadri, P.; Chen, H.; Singh, S.; Goldfarb-Tarrant, S. Small Changes, Large Consequences: Analyzing the Allocational Fairness of LLMs in Hiring Contexts. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*; Inui, K., Sakti, S., Wang, H., Wong, D.F., Bhattacharyya, P., Banerjee, B., Ekbal, A., Chakraborty, T., Singh, D.P., Eds.; The Asian Federation of Natural Language Processing and The Association for Computational Linguistics: Mumbai, India, 2025; pp. 2645–2665. [CrossRef]
18. Prasad, K.D.V.; De, T. Generative AI as a catalyst for HRM practices: Mediating effects of trust. *Humanit. Soc. Sci. Commun.* **2024**, *11*, 1362. [CrossRef]
19. Tsiskaridze, R.; Reinhold, K.; Jarvis, M. Innovating HRM Recruitment: A Comprehensive Review of AI Deployment. *Mark. Manag. Innov.* **2023**, *14*, 239–254. [CrossRef]
20. Aleisa, M.; Alshahrani, M.; Beloff, N.; White, M. TAIRA-BSC—Trusting AI in Recruitment Applications through Blockchain Smart Contracts. In Proceedings of the 2022 5th IEEE International Conference on Blockchain (Blockchain 2022), Espoo, Finland, 22–25 August 2022; pp. 376–383. [CrossRef]
21. Kinger, S.; Kinger, D.; Thakkar, S.; Bhake, D. Towards smarter hiring: Resume parsing and ranking with YOLOv5 and DistilBERT. *Multimed. Tools Appl.* **2024**, *83*, 82069–82087. [CrossRef]
22. Frissen, R.; Adebayo, K.; Nanda, R. A Machine Learning Approach to Recognize Bias and Discrimination in Job Advertisements. *AI Soc.* **2023**, *38*, 1025–1038. [CrossRef]
23. Senger, E.; Zhang, M.; van der Goot, R.; Plank, B. Deep Learning-based Computational Job Market Analysis: A Survey on Skill Extraction and Classification from Job Postings. In *Proceedings of the First Workshop on Natural Language Processing for Human Resources (NLP4HR 2024)*; Hruschka, E., Lake, T., Otani, N., Mitchell, T., Eds.; Association for Computational Linguistics: St. Julian's, Malta, 2024; pp. 1–15. [CrossRef]
24. Zhang, M.; Jensen, K.; Sonniks, S.; Plank, B. SkillSpan: Hard and Soft Skill Extraction from English Job Postings. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*; Carpuat, M., de Marneffe, M.C., Meza Ruiz, I.V., Eds.; Association for Computational Linguistics: Seattle, WA, USA, 2022; pp. 4962–4984. [CrossRef]
25. Urbano, J.; Couto, M.; Rocha, G.; Cardoso, H.L. Inconsistency Detection in Job Postings. In *Proceedings of the 3rd Conference on Language, Data and Knowledge (LDK 2021)*; Open Access Series in Informatics (OASIcs); Schloss Dagstuhl—Leibniz-Zentrum für Informatik: Wadern, Germany, 2021; Volume 93, pp. 25:1–25:16. [CrossRef]
26. Arnold, D.; Quach, S.; Taska, B. *The Impact of Pay Transparency in Job Postings on the Labor Market*; Working Paper 34480; National Bureau of Economic Research: Cambridge, MA, USA, 2025. Available online: https://www.nber.org/papers/w34480 (accessed on 16 April 2026).
27. Audoly, R.; Bhuller, M.; Reiremo, T.A. *The Pay and Non-Pay Content of Job Ads*; Staff Reports 1124; Federal Reserve Bank of New York: New York, NY, USA, 2024. [CrossRef]
28. Naudé, M.; Adebayo, K.; Nanda, R. A Machine Learning Approach to Detecting Fraudulent Job Types. *AI Soc.* **2023**, *38*, 1013–1024. [CrossRef]
29. Năstase, M.; Croitoru, G.; Florea, N.V.; Cristache, N.; Lile, R. The Perceptions of Employees from Romanian Companies on Adoption of Artificial Intelligence in Recruitment and Selection Processes. *Amfiteatru Econ.* **2024**, *26*, 421–439. [CrossRef]

30. Schloetzer, J.D.; Yoshinaga, K. Algorithmic Hiring Systems: Implications and Recommendations for Organisations and Policymakers. In *YSEC Yearbook of Socio-Economic Constitutions 2023: Law and the Governance of Artificial Intelligence*; Gill-Pedro, E., Moberg, A., Eds.; Springer: Cham, Switzerland, 2024; pp. 213–246. [CrossRef]
31. Peffers, K.; Tuunanen, T.; Gengler, C.E.; Rossi, M.; Hui, W.; Virtanen, V.; Bragge, J. Design Science Research Process: A Model for Producing and Presenting Information Systems Research. *arXiv* **2020**, arXiv:2006.02763. [CrossRef]
32. Ferreira Cruz, E.; Rosado da Cruz, A.M. Design Science Research for IS/IT Projects: Focus on Digital Transformation. In *Proceedings of the 15th Iberian Conference on Information Systems and Technologies (CISTI 2020)*; IEEE: New York, NY, USA, 2020; pp. 1–6. [CrossRef]
33. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design Science in Information Systems Research. *MIS Q.* **2004**, *28*, 75–105. [CrossRef]
34. Mahjoub, A.; Kruyen, P.M. Efficient recruitment with effective job advertisement: An exploratory literature review and research agenda. *Int. J. Organ. Theory Behav.* **2021**, *24*, 107–125. [CrossRef]
35. Touvron, H.; Lavril, M.; Izacard, G.; Martinet, X.; Lachaux, M.A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. LLaMA: Open and Efficient Foundation Language Models. *arXiv* **2023**, arXiv:2302.13971. [CrossRef]
36. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*; ACM: New York, NY, USA, 2017; pp. 6000–6010. Available online: https://dl.acm.org/doi/10.5555/3295222.3295349 (accessed on 16 April 2026).
37. Boonstra, L. *Prompt Engineering*; Technical Report; Google: Mountain View, CA, USA, 2025.
38. Paradis, T. Entry-Level Jobs in Areas Like Tech Often Require Years of Experience. Business Insider. Available online: https://www.businessinsider.com/entry-level-jobs-tech-roles-require-years-of-experience-2024-10 (accessed on 12 December 2025).
39. Mohr, T.S. Why Women Don't Apply for Jobs Unless They're 100% Qualified. Harvard Business Review. Available online: https://hbr.org/2014/08/why-women-dont-apply-for-jobs-unless-theyre-100-qualified (accessed on 12 December 2025).
40. Johnson, D.; Menezes, A.; Vanstone, S. The Elliptic Curve Digital Signature Algorithm (ECDSA). *Int. J. Inf. Secur.* **2001**, *1*, 36–63. [CrossRef]
41. Liu, J. Digital signature and hash algorithms used in Bitcoin and Ethereum. In *Proceedings of the Third International Conference on Machine Learning and Computer Application (ICMLCA 2022)*; Ba, S., Zhou, F., Eds.; International Society for Optics and Photonics, SPIE: Bellingham, WA, USA, 2023; Volume 12636, p. 126365H. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.