Humanizing the ATS-Based Recruiting Process by Using LLMs and Putting the Human Back in the Loop

Valdo V. Mpinga * and A. M. Rosado da Cruz *

Posted Date: 17 March 2026

doi: 10.20944/preprints202603.1274.v1

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)

 doi:10.20944/preprints202603.1274.v1

Article

Humanizing the ATS-Based Recruiting Process by Using LLMs and Putting the Human Back in the Loop

Valdo V. Mpinga [1,*] and A. M. Rosado da Cruz [2,*]

[1] 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 (ATS) have evolved significantly since their inception in 1996, transitioning from simple *resumé* repositories to AI-driven tools with advanced capabilities. Although these innovations have improved recruitment efficiency, they have also introduced significant ethical and Human Rights challenges. Bias in machine learning (ML) training data and over reliance on algorithmic decision making have led to violations of Human dignity, equality, privacy, and other fundamental rights. This study examines these challenges and proposes a Human-centered approach to recruitment processes that integrates ATS with ethical safeguards. The proposed methodology involves the development of a Humanization application service to mitigate bias and enhance Human oversight. Key features include validating job vacancy requirements, identifying and addressing bias triggers in recruitment algorithms, and requiring digital signatures from qualified professionals to approve job postings, ensuring that there are humans that assume responsibility. By incorporating generative AI and blockchain, this system facilitates compliance, transparency, and fairness in recruitment, thereby addressing the ethical concerns of algorithmic hiring. Crucially, this study demonstrates that organizations can achieve meaningful Humanization of their recruitment processes, including the mitigation of bias and provision of personalized feedback (among many other potential steps) through targeted interventions that require surprisingly modest financial investment, time commitment, and computational resources. This study underscores the importance of combining technological advancements with ethical considerations to create equitable recruitment processes in an accessible and resource-efficient manner.

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 (ATS) are a type of software used by companies, typically by management and human resources (HR) teams, to efficiently manage job openings, and respective job advertisements, and the applications for those openings. In 1996, the first ATS systems appeared, with features for storing applicants *resumés*, keyword filtering and searching, advertise job openings, and candidate tracking [1]. Years later, with the advancement of technology, this type of software evolved, being able to do much more. The rapid development of Artificial Intelligence (AI), especially Machine Learning (ML), and its integration into ATS systems have enabled an increase in productivity in HR management teams. Due to the AI/ML integration into ATS 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 certain dependence and accommodation in HR management teams due to the practice of recruitment based mainly on algorithmic management.

 doi:10.20944/preprints202603.1274.v1

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 ATS systems, 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 ATS systems 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 on how the system is used and the bias on the dataset used for training the ML models.

With the goal of preventing discrimination, and driven by the lack of regulation in AI-driven recruitment, the authors in [6] have developed a Multi-Agent System architecture to facilitate ethical and legal auditing of AI-driven recruitment, emphasizing the analysis of job video interviews. Another major issue arises when the recruitment process is driven solely by the algorithmic management.

 doi:10.20944/preprints202603.1274.v1

When an human actor is not involved, purely algorithmic recruitment processes tend to be perceived as unfair by job applicants [7]. Paramita *et al.* revealed that both transactional and relational aspects of AI in recruitment are equally important in delivering effective recruitment services [8]. When these factors are ignored, what we often obtain is human dignity violation by unfair algorithm assessments, as revealed by [9]. These assessments impose a reductionist, fixed view of candidates and infringe on their human dignity. Unfortunately, human dignity and human rights are violated, as hiring algorithms can inflict significant harm to individual human rights, including the right to work, equality and non-discrimination, privacy, freedom of expression and of association [2].

Another major problem in today's hiring algorithms is bias. In [10], Fabris et al. examined the role of bias-conducive factors in algorithmic hiring, highlighting the importance of considering group fairness, granularity, and interpretability in fairness measures, and discussed various mitigation strategies to address discrimination in the hiring process. As a result, the study confirmed that various bias-conducive factors such as age, disability, gender, religion or belief, racial or ethnic origin, and sexual orientation can lead to unfair recruitment systems. The authors also highlight the importance of considering sensitive attributes in algorithmic hiring, and note that there is a lack of consensus on how to treat algorithms that favor historically disadvantaged communities [10]. 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 [11], 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 [11].

Other studies have focused on solving the problems highlighted in previous studies. Tsiskaridze *et al.* studied innovation in the Human Resources field and found that there is growing academic interest in AI utilization in HRM, motivated by the potential of AI to mitigate personal biases and provide more trustworthy information, and the importance of balancing human and AI elements in HRM [12]. Additionally, the study highlights the need for transparency, fairness, and ethical considerations in AI-based decision making. In [13], the potential of integrating AI and Blockchain Technology (BCT) is highlighted, to improve transparency and trust in the recruitment process. Aleisa *et al.* proposed a conceptual architecture for a proof of concept tool that can serve as a foundation for future studies focused on improving trust in AI applications through the integration of BCT. Another important factor in an applicant evaluation is the input, which comes from the applicant's *resumé*. So if the tool extracting the data from it performs poorly, it could be very misleading. In [14], Kinger et al. proposed a method that integrates cutting-edge technologies to refine the existing ATS process, offering a tailored and precise approach to evaluating *resumés*. The system achieved an accuracy of 96.2% in *resumé* parsing, thereby significantly improving the efficiency of the candidates selection process.

Nastase *et al.* evaluated the perception of employees of a company, and the results showed that almost all the variables proposed for the model positively influenced the intention to accept and use AI in the recruitment and selection process [15]. Non-discrimination and the use of AI in the selection and recruitment processes had little influence.

In summary, although significant progress has been made in understanding the ethical, legal, and technical implications of AI-driven ATS systems, several issues remain unresolved. Ethical auditing frameworks such as those proposed in [6,16], are largely conceptual and lack empirical validation with real-world data. Also, studies emphasizing the importance of balancing algorithmic and human judgment in recruitment [8,9] often do not specify how such a balance should be operationalized. Research is still needed to determine the optimal degree of automation and the appropriate mechanisms for human oversight to safeguard fairness and human dignity. There is also no clear consensus on how to define or measure fairness or how to handle sensitive attributes in mitigating bias [10]. Fabris et al. (2024) highlight that fairness metrics vary depending on whether sensitive attributes are used or

 doi:10.20944/preprints202603.1274.v1

excluded and that the treatment of algorithms favoring historically disadvantaged groups remains controversial [10].

3. Methodology

The methodology followed in this research is based on Design Science Research (DSR) [17,18]. DSR is well suited to this study because it focuses on creating and evaluating artifacts that solve identified real-world problems while generating both practical and theoretical contributions [17,19].

Having identified real-world, socio-technical problems (ethical, technical, and organizational ATS issues), the goal is to design and develop innovative artifacts, and evaluate their effectiveness through rigorous testing and refinement, to improve ethical and organizational outcomes in ATS systems. This approach ensures that technological advancements remain aligned with moral values, regulatory requirements, and practical human resource practices. To structure this study, the six activities of the DSR process were followed [17,18] (see Figure 1).

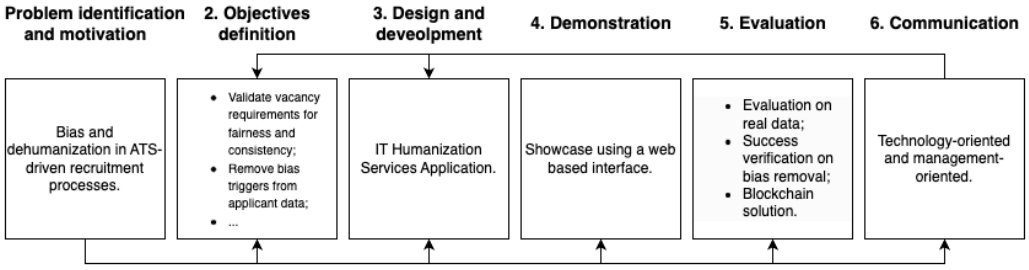

Figure 1. DSR diagram applied to the project's context.

3.1. Problem Identification and Motivation

This study addresses the critical problems identified before, namely:

- The potential for algorithmic bias in screening candidates, leading to unfair disadvantages for certain groups. This is derived from the lack of consensus on how fairness should be defined, measured, or operationalized in AI-based hiring. Particularly, the treatment of sensitive attributes such as gender, race, and age may be a bias-conducive factor or bias trigger [2,10,20].
- The lack of transparency in how AI evaluates candidates and the limited empirical validation or practical implementation in real-world contexts, hinders trust in the hiring process [6,7,9].
- Keeping a balance between human and algorithmic decision-making, ensuring the appropriate distribution of decision-making between humans and AI systems. Keeping the human in the loop, in recruitment processes, is essential to avoid overreliance on automated judgments and ensure ethical considerations are paramount when integrating AI into recruitment [8,9]

Addressing these problems is significant as it promotes more equitable hiring practices, enhances organizational reputation, and aligns technological advancements with societal values. This aligns with the DSR principle of developing technology-based solutions for relevant business problems.

3.2. Objectives Definition

The objectives of this research are to design and develop a humanization service application that:

- Validates job opening's requirements to ensure fairness and consistency in job postings. This involves analyzing job descriptions to identify and rectify misalignments, unreasonable demands, unrealistic expectations, and internal discrepancies.
- Removes bias triggers from applicant data to mitigate the risk of algorithmic bias. This focuses on identifying and removing sensitive information that could lead to unfair evaluations, while preserving the integrity of the original data.

 doi:10.20944/preprints202603.1274.v1

- Implements a digital signature mechanism for human reviewers to enhance accountability and transparency. This involves a decentralized validation protocol using blockchain to ensure that both HR personnel and subject-matter experts authorize job openings publications.

These objectives are qualitative, focusing on how the application of Humanization Services improves the fairness and transparency of recruitment processes. They are designed to create disruptions that trigger strategic responses from organizations.

3.3. Design and Development

This activity involved the design and construction of the humanization services. Key components include:

- **Vacancy Requirement Validation Module:** Developed using Python and the Fast API, incorporating LLaMA prompts to evaluate job postings. It analyzes job descriptions and provides structured feedback regarding fairness and consistency.
- **Bias Triggers Removal Module:** Implemented with Python, the Fast API and LLaMA prompts. This module identifies and removes bias-related fields from applicant's data, ensuring the original data format is maintained.
- **Digital Signature of Relevant Human Actors Module:** A decentralized validation protocol was designed, leveraging Blockchain technology, to mandate dual cryptographic authorization from HR personnel and subject-matter experts.

These artifacts are the practical outputs of our development research, including models, methods, and new technical features.

3.4. Demonstration

The demonstration of the humanization service application showcases its functionality in a real-world or simulated environment. Key demonstrations included:

- Validating a set of job postings to illustrate the application's ability to identify inconsistencies and biases.
- Processing applicant's data to demonstrate the removal of bias triggers, while preserving data integrity.
- Simulating the blockchain-based validation process to confirm the correct implementation of access control and signature verification.
- A web interface has been developed to showcase job analysis, bias-free candidate view, and blockchain integration (available at https://joblimpo.valdompinga.com/).

These demonstrations illustrate the use of artifacts to solve instances of the identified problems.

3.5. Evaluation

The evaluation phase assessed the effectiveness of the application of Humanization Services in achieving defined objectives. The evaluations included:

- Analyzing the results of job openings' validation on a dataset of 21,701 job postings, measuring the percentage of postings with misalignments, unreasonable demands, unrealistic expectations and internal discrepancies.
- Verifying the successful removal of bias-related fields from applicant's data.
- Testing the blockchain-based validation protocol, to ensure correct access control, signature verification, and workflow integrity.

The evaluation involved comparing the results with the objectives defined and employing relevant analysis techniques and metrics.

 doi:10.20944/preprints202603.1274.v1

3.6. Communication

The findings of this research are communicated in this article, highlighting the relevance of the problem, the design and utility of the Humanization Services application, and the results of the evaluation. Communication is tailored to both technology- and management-oriented audiences.

- For technology-oriented audiences (e.g., AI developers, software engineers), detailed information is provided on the application's architecture, implementation using Python, Fast API and LLaMA, and the blockchain protocol.
- For management-oriented audiences (e.g., HR managers, organizational leaders), the focus is on the application's ability to enhance fairness and transparency, mitigate legal risks, and improve organizational reputation.

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.

 doi:10.20944/preprints202603.1274.v1

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.	Applicants data or resumé (in JSON, XML or plain text format)	Applicants 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 of 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 requirements 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 efficient talent acquisition arises from the prevalence of incongruous job specifications. A substantial portion of these issues could be mitigated through the implementation of a validation mechanism involving human subject matter experts with a profound understanding of the requisite technical skills. This additional layer of scrutiny serves not only to prevent the inadvertent exclusion of suitably qualified candidates, but also reinforces the accuracy and relevance of the initial requirements definition. To facilitate this crucial validation step, a graphical user interface is implemented, enabling at least one individual with pertinent field expertise, such as a project manager, to review and explicitly approve the vacancy requirements prior to public dissemination, subsequent to the automated humanization process. The underlying technology employed to ensure the integrity and traceability of this approval workflow is a distributed ledger system based on the blockchain principles.

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 validate and subscribe a job specification (refer to Table 1). Interaction is initiated when the ATS transmits either a vacancy description or applicant's *résumés*

 doi:10.20944/preprints202603.1274.v1

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 services API.

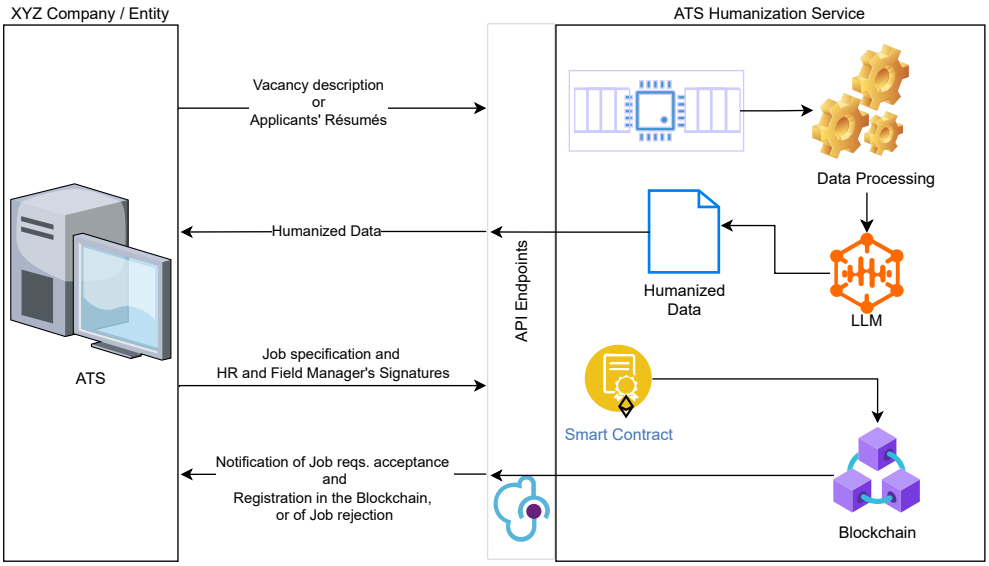

Figure 2. Solution architecture.

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 state-of-the-art natural language processing (NLP) model designed to advance research in generative AI [21]. It is a transformer-based model trained on an extensive and diverse corpus of text, offering high performance with relatively fewer parameters than other large-scale models, such as GPT-4, while achieving competitive results across various NLP benchmarks [21].

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 [22]. 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

 doi:10.20944/preprints202603.1274.v1

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 requirements 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 has been developed and is accessible at the URL https://joblimpo.valdompinga.com/. 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 **Landing Page** serves as the initial point of contact for the users (https://joblimpo.valdompinga.com/). 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, allows any user to analyze a job description against predefined ethical and practical criteria. Users can input a job description or requirements' text 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, 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, 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 [23] 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 Validater. 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.*

 doi:10.20944/preprints202603.1274.v1

- ***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*

 doi:10.20944/preprints202603.1274.v1

- *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 encounter challenges related to the inclusion of unrealistic requirements, often stemming from a lack of sufficient domain-specific knowledge during the drafting phase [24,25].

To mitigate this critical issue, we advocate that the final decision and responsibility for a job posting should rest with a human. For ensuring this, we propose a decentralized validation protocol based on the necessity of dual cryptographic authorization from both HR personnel and subject-matter experts, 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.

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 [26,27].
- **Immutable Approval Records**: Secure and transparent recording of all approval processes through the immutable state transitions inherent to the blockchain.

 doi:10.20944/preprints202603.1274.v1

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/), or private/protected, as is the case of Hyperledger Besu (https://besu.hyperledger.org/).

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.

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

 doi:10.20944/preprints202603.1274.v1

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.

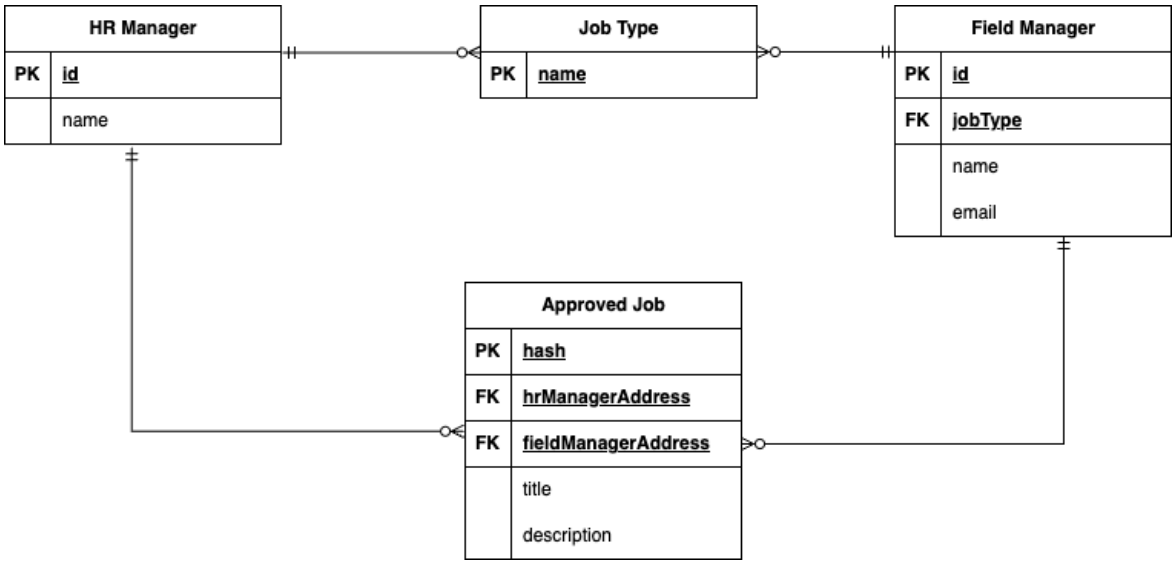

Figure 3. Data model of the Hardhat Job Approval smart contract.

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.

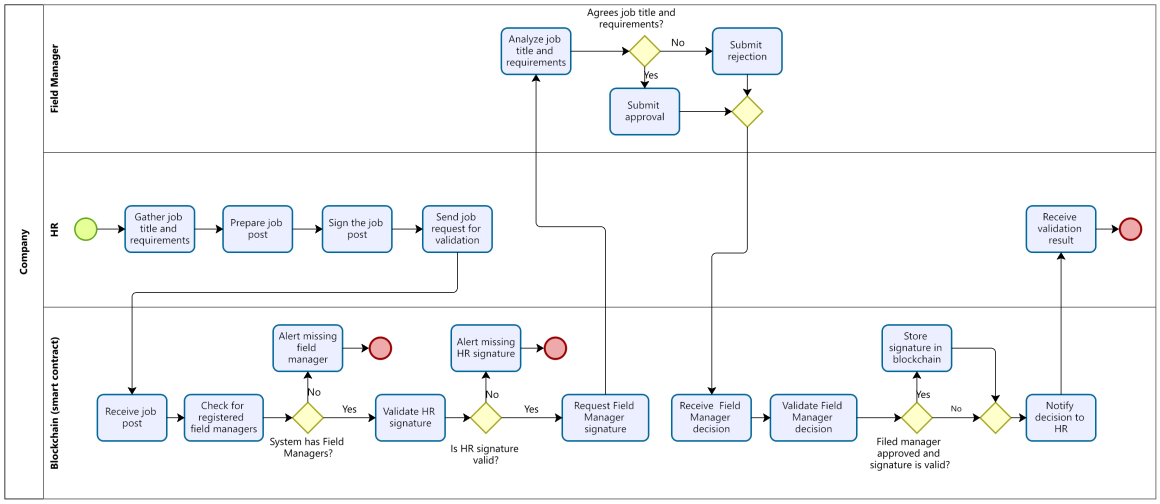

Figure 4. Detailed job posting validation workflow utilizing dual cryptographic signatures.

 doi:10.20944/preprints202603.1274.v1

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 deploy, the HR manager is granted administrative privileges that enable them to:

- **HR Personnel Onboarding**: Enroll additional HR personnel into the system using the `addHRManager()` function of the smart contract.
- **Domain Expert Registration**: Register qualified domain experts within the system, associating them with specific job-type taxonomies, using the smart contract function `addFieldManager(jobType)`.
- **Job-Type Taxonomy Configuration**: Define and manage the various categories of job types recognized by the system (e.g., "Software Engineering", "Biomedical Engineering").

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

 doi:10.20944/preprints202603.1274.v1

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 hours. 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 a notable achievement 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):

- 64.76%, with a 95% confidence interval (CI) between 64.12%-–65.39%, out of the analyzed 21,701 job postings failed our Role Alignment check (14,053 out of 21,701).
- 35.99% of the jobs, with 95% CI between 35.35%–36.63%, had unreasonable experience requirements (7,816 out of 21,701 failed Experience Rationality).
- 66.56% of the workloads, with 95% CI between 65.93%–67.19%, were unrealistic for the role of one person (14,445 out of 21,701 failed Workload Feasibility).
- 38.11% of the jobs, with 95% CI between 37.46%–38.76%, contained discrepancies (8,270 out of 21,701).

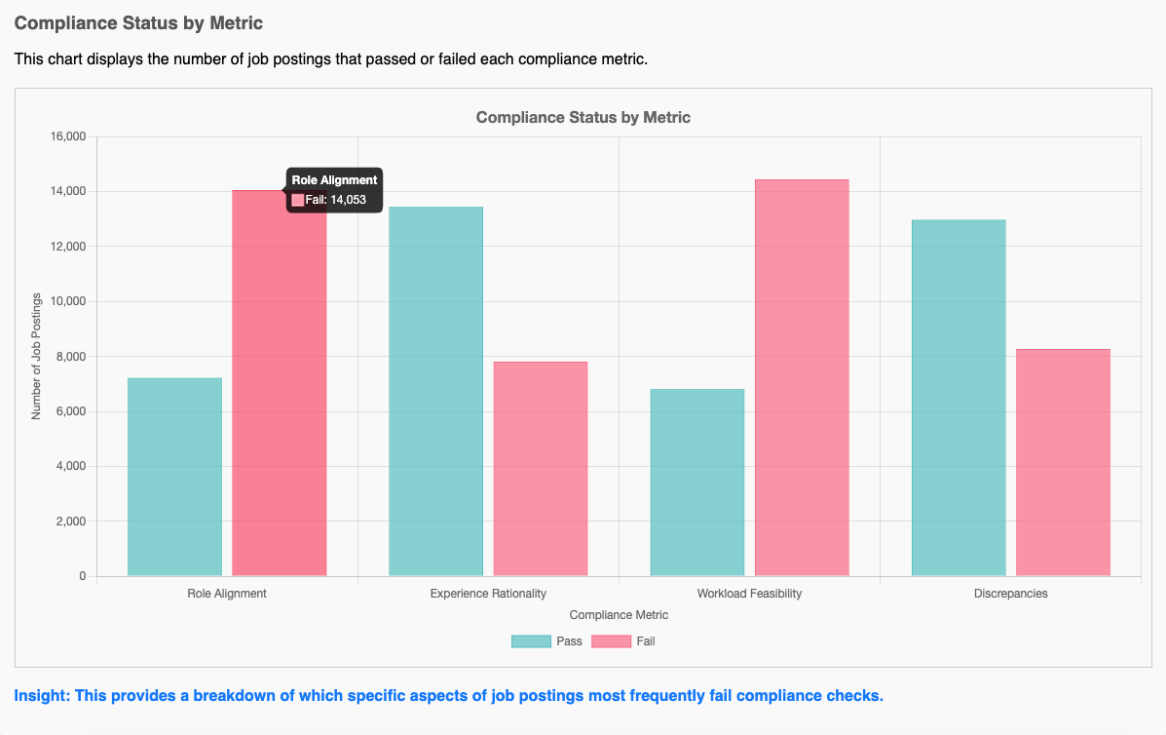

Figure 5. Evaluation stats for the vacancy validator service.

5.2. Revealing Bias Triggers

The service endpoint was determined as expected. Typically, ATS systems 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 following format.

```
{
```

 doi:10.20944/preprints202603.1274.v1


```
  "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"
}
```


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 endpoint successfully produces this output for JSON, XML, and plain-text data.

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. This evaluation aimed to assess their capability to produce structured, pattern-like output, a crucial aspect for automating the analysis of multiple job requirements. By comparing models of different sizes, we sought to understand the trade-offs between computational efficiency and analytical accuracy in this specific application.

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 [21]. 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 requirements analysis, we selected a range of open-source models categorized into three size groups to represent different computational resource constraints (See Table 3).

 doi:10.20944/preprints202603.1274.v1

Table 3. Evaluation of Large Language Models for Job Requirements 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	...
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

 doi:10.20944/preprints202603.1274.v1

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 0.27s to 2.37s, Llama-3.2-1B from 1.86s to 4.15s) exhibited the fastest processing times, often below 5 seconds per job posting. The Light models showed increased variability, with Phi-3-Mini ranging from 3.84s to 14.01s, while the Medium-Light models (e.g., Mistral-7B from 5.97s to 13.31s, Llama-3.1-8B from 5.50s to 17.86s) generally took longer, with some exceeding 10 seconds. 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. **Workload Feasibility** and **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 40-80 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

 doi:10.20944/preprints202603.1274.v1

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**: The tests demonstrated a 100% detection rate for both invalid HR and field manager digital signatures, ensuring that only cryptographically authorized personnel can endorse job postings (See Figure 6).
- **Unauthorized Approval Prevention**: The ECDSA (Elliptic Curve Digital Signature Algorithm) validation mechanism effectively prevented all attempts of unauthorized job approvals by entities lacking the required valid digital signatures.
- **Duplicate Submission Blocking**: The implemented job hashing mechanism successfully blocked all attempts to submit and approve duplicate job postings, maintaining the integrity and uniqueness of validated vacancies.

 doi:10.20944/preprints202603.1274.v1


```
Compiled 2 Solidity files successfully (evm target: paris).


  JobApproval
    HR Manager CRUD
      ✔ Should initialize with one HR manager
      ✔ Should add a new HR manager
      ✔ Should not allow non-HR to add HR managers
      ✔ Should update an HR manager's name
      ✔ Should remove an HR manager
    Field Manager CRUD
      ✔ Should add a field manager
      ✔ Should not allow non-HR to add field managers
      ✔ Should update a field manager's details
      ✔ Should remove a field manager
    Job Approval
      ✔ Should allow HR and Field Manager to sign a job approval
      ✔ Should fail if HR signature is incorrect
      ✔ Should fail if Field Manager signature is incorrect
      ✔ Should fail if no field manager is set for the job type
      ✔ Should not allow duplicate job approvals
      ✔ Should store and retrieve job approvals correctly

  Lock
    Deployment
      ✔ Should set the right unlockTime
      ✔ Should set the right owner
      ✔ Should receive and store the funds to lock
      ✔ Should fail if the unlockTime is not in the future
    Withdrawals
      Validations
        ✔ Should revert with the right error if called too soon
        ✔ Should revert with the right error if called from another account
        ✔ Shouldn't fail if the unlockTime has arrived and the owner calls it
      Events
        ✔ Should emit an event on withdrawals
      Transfers
        ✔ Should transfer the funds to the owner


  24 passing (547ms)

○ valdo@Valdos-MacBook-Pro clean-job %
```

Figure 6. 100% test coverage.

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.

 doi:10.20944/preprints202603.1274.v1

- 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.

5.5. System Demonstration via Public Web Interface

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

5.5.1. Interactive Job Requirement Analysis

- **Accessibility**: The job requirement analysis tool is publicly accessible via the following URL: https://joblimpo.valdompinga.com/requirements.
- **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.

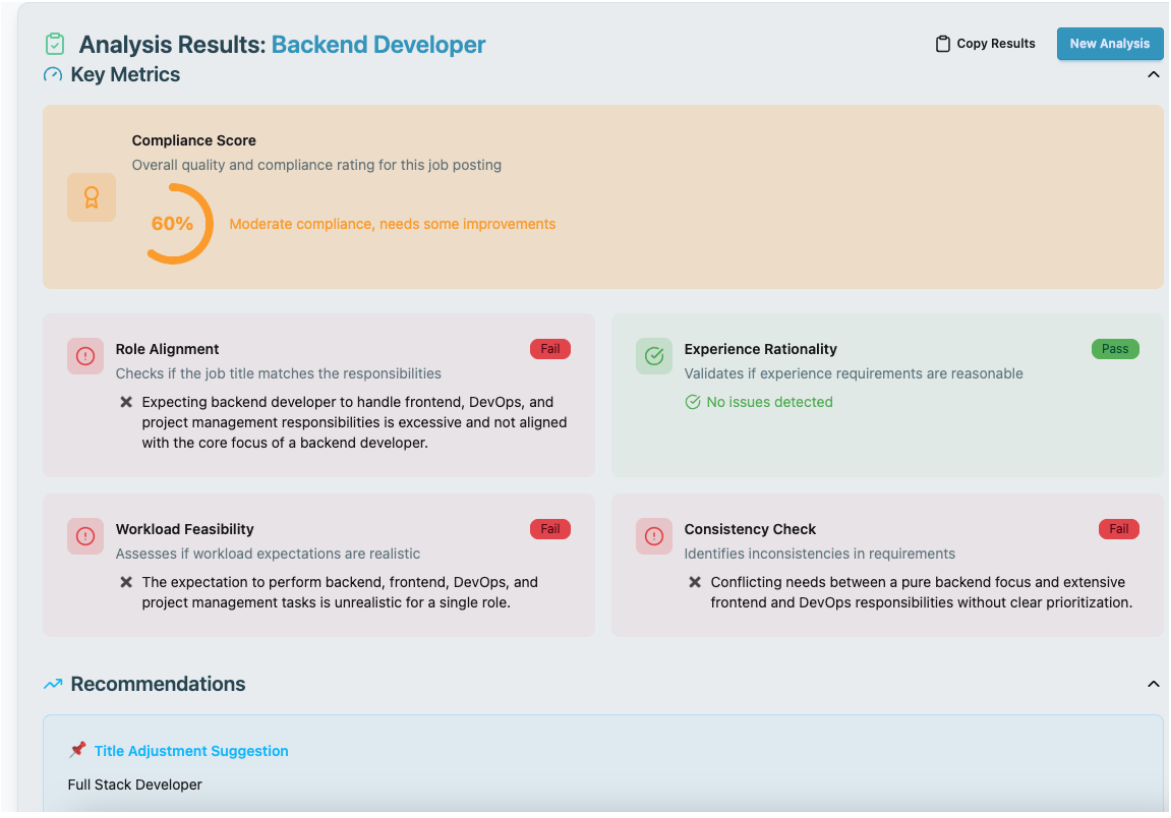

Figure 7. Interactive web interface for job requirement evaluation, displaying dehumanization detection metrics derived from natural language analysis.

5.5.2. Bias-Free Candidate Presentation Interface

- **Accessibility**: The candidate anonymization service demonstration is available at: https://joblimpo.valdompinga.com/candidate.
- **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.

 doi:10.20944/preprints202603.1274.v1

- **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.

JSON Format.

Figure 8 illustrates the anonymization process for candidate data provided in the JSON format.

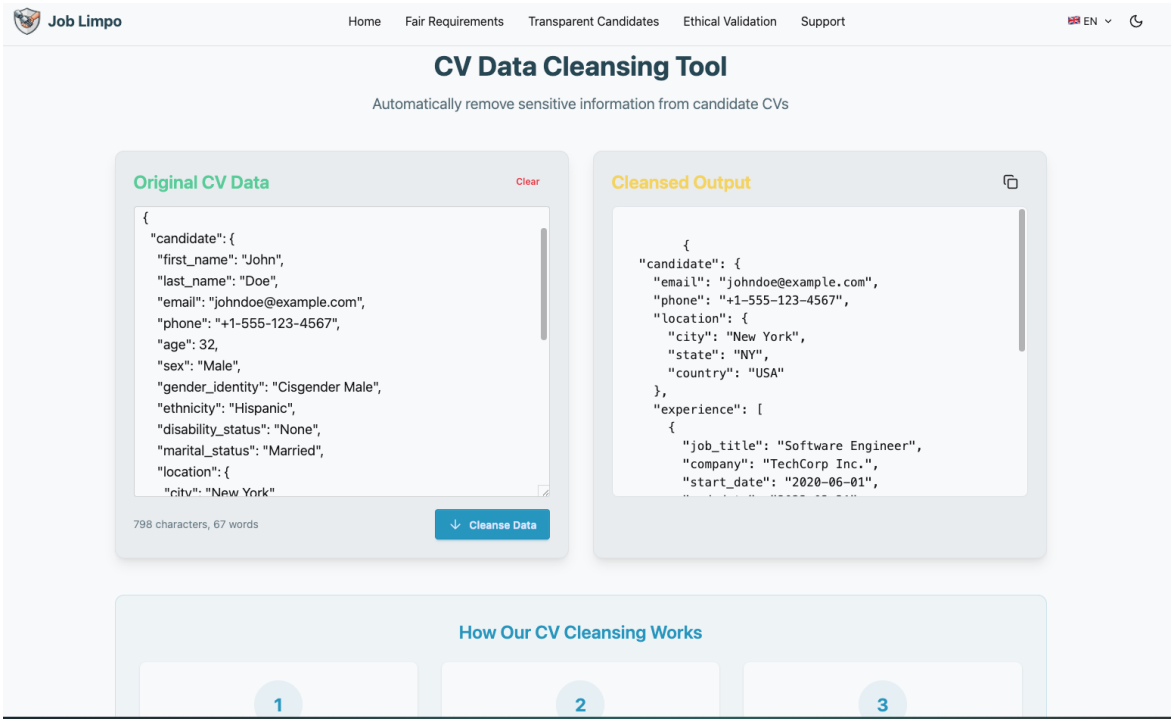

Figure 8. Web interface demonstrating candidate profile processing with bias-inducing information removal for JSON format.

XML Format.

Figure 9 illustrates the anonymization process for candidate data provided in the XML format.

 doi:10.20944/preprints202603.1274.v1

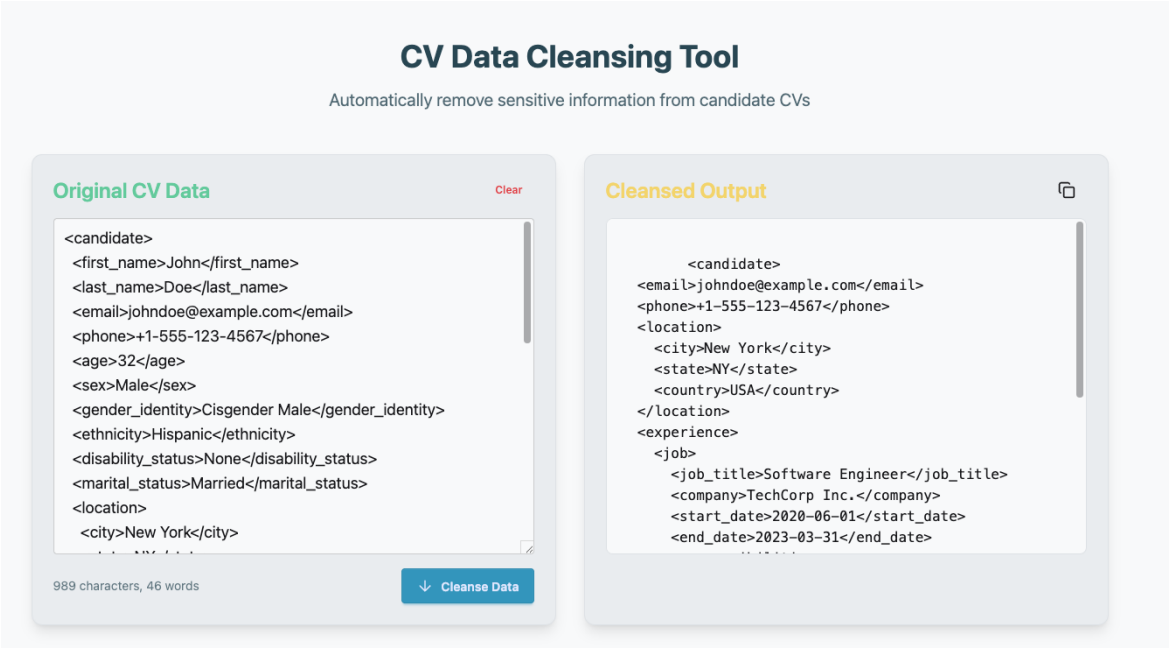

Figure 9. Web interface demonstrating candidate profile processing with bias-inducing information removal for XML format.

Plain Text Format.

Figure 10 illustrates the anonymization process for the candidate data provided in plain-text format.

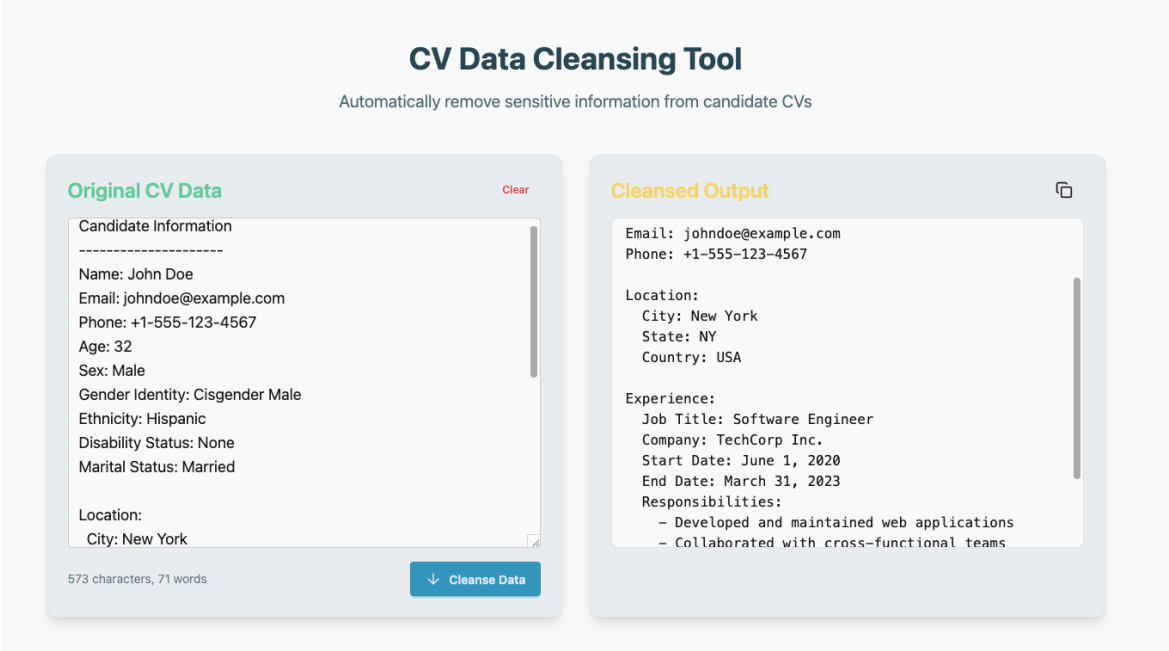

Figure 10. Web interface demonstrating candidate profile processing with bias-inducing information removal for plain text format.

5.5.3. Blockchain-Based Service Implementation Showcase

- **Platform Access**: The interface providing access to the blockchain implementation details is hosted at: https://joblimpo.valdompinga.com/validation.
- **Open-Source Repository Link**: This section provides a direct link to the project's public GitHub repository, which contains the following critical components:

 doi:10.20944/preprints202603.1274.v1

- 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.

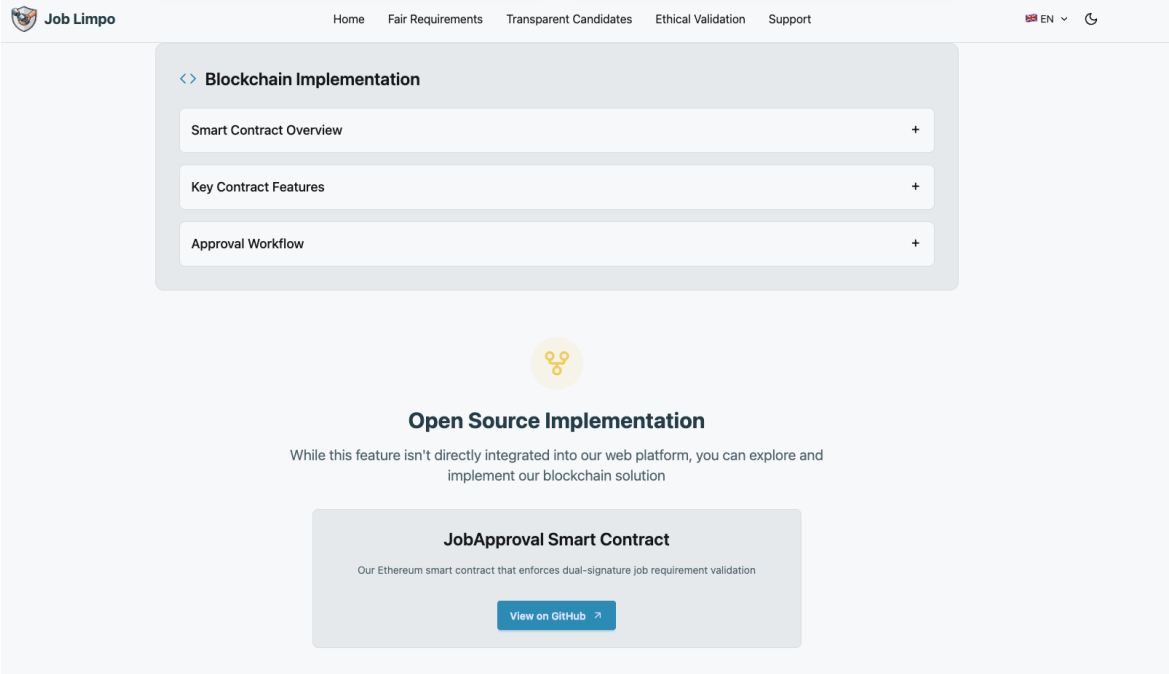

Figure 11. Web interface providing a link to the open-source blockchain solution repository.

5.6. Limitations and Future Work

At this point, several study limitations must be acknowledged. Despite having presented a demonstration of the designed and implemented platform, a user validation study, or user acceptance tests, have not 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. Conclusion

The persistent dehumanization of job roles and the inadvertent violation of fundamental human rights within recruitment practices constitute a significant and ethical concern. Empirical evidence robustly demonstrates that neglecting these critical issues not only undermines the inherent dignity of individuals but also directly contravenes the established tenets of human rights law. The accelerating proliferation of AI across various societal domains, including talent acquisition, has led to the increasing deployment of sophisticated AI algorithms within ATS and broader recruitment workflows. Alarmingly, these algorithms frequently exhibit a propensity to inherit and amplify preexisting biases present in their training data. This bias amplification typically occurs unintentionally through data collection and model training processes that inadvertently embed historical societal inequities.

To cultivate a more inclusive talent acquisition environment capable of attracting a diverse pool of highly qualified candidates, the presence of rational and equitable job vacancies within the labor market

 doi:10.20944/preprints202603.1274.v1

is indispensable. Regrettably, this ideal is not always realized in practice. Companies often construct job postings based on internally perceived "appropriate" keywords, sometimes leading to the inclusion of unrealistic or overly specific requirements, and occasionally even stipulating levels of prior experience that exceed pragmatic industry norms. The central tenet of this study was to provide compelling evidence that organizations can effectively humanize the multifaceted aspects of their recruitment processes without incurring exorbitant expenditures of financial capital, extensive time commitments, or the need for vast computational infrastructure. By focusing on targeted interventions and leveraging accessible technologies, this study demonstrates a pathway towards more ethical and equitable talent acquisition. For improving balance between the efficiency offered by AI and ATS systems, and the essential need for human judgment to ensure fair and unbiased hiring decisions, three validation endpoints have been proposed: a Job Requirements Validation Module; a Bias Triggers Removal Module; and a blockchain-based decentralized validation module that ensures the involvement and accountability of a HR manager and a human field expert.

The imperative to humanize recruitment extends beyond mitigating bias in job requirements and user data. Consider the often-neglected aspect of candidate feedback. It is demonstrably inefficient and arguably dehumanizing to expect human resource professionals to provide personalized feedback to thousands of applicants at each vacancy. LLMs offer scalable and cost-effective solutions. By leveraging LLM, organizations can generate tailored feedback based on the specific requirements of the roles and attributes of individual candidates. This not only enhances the candidate experience but also fosters a sense of respect and value, transforming a traditional impersonal process into a more human-centric interaction.

The practical implementation of such solutions is surprisingly accessible even for enterprises with budgetary constraints. First, numerous powerful LLM models are available under open-source licenses, significantly reducing software acquisition costs. Secondly, the technical barrier for deploying these models is relatively low, often requiring only basic programming skills and not an extensive computational infrastructure. Standard enterprise-grade hardware is frequently used in such applications. Consequently, the integration of LLMs for tasks such as personalized feedback represents a tangible and affordable step for organizations to actively contribute to a more ethical and equitable recruitment landscape, shifting away from practices that inadvertently violate job seekers' fundamental rights and dignity. Using a blockchain for the digital dual-signature of job requirements ensures human involvement and accountability (HR manager and Field-manager) in validating each job post, before its deployment, in an immutable and tamper-proof registry.

Besides the above mentioned opportunities to humanize recruitment, future work should address the limitations identified before. It should include empirical validation and user acceptance testing with HR professionals to assess how effectively fairness and job specifications/ job function alignment is improved by the proposed approach. The lack of comparable studies or published humanization approaches reporting similar metrics will demand the development of suitable benchmark datasets for comparative analysis.

Since we are dealing with sensitive personal data, a deeper exploration of how the proposed system aligns with GDPR and other private-data-protection law frameworks is also something that should be addressed in future work.

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.; writing—review and editing, V.V.M. and 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 National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia within project: UID/06121/2023.

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

 doi:10.20944/preprints202603.1274.v1

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

References

1. Recruitify.ai. The Evolution of Applicant Tracking Systems (ATS): From Manual Processes to AI-powered Recruitment. https://www.recruitify.ai/blog/en/the-evolution-of-applicant-tracking-systems-ats-from-manual-processes-to-ai-powered-recruitment, 2023. Accessed: 2025-05-01.
2. Yam, J.; Skorburg, J.A. From human resources to human rights: Impact assessments for hiring algorithms. *Ethics and Information Technology* **2021**, *23*, 611 – 623. https://doi.org/10.1007/s10676-021-09599-7.
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.
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, Forty-fifth session, 26 February–5 April 2024.
5. Amnesty International. Surveillance Giants: How The Business Model of Google and Facebook Threatens Human Rights, 2019. (Accessed: 2025-07-23).
6. Del Carmen Fernández Martínez, M.; Fernández, A. AI in recruiting. Multi-agent systems architecture for ethical and legal auditing. *IJCAI International Joint Conference on Artificial Intelligence* **2019**, *2019-August*, 6428 – 6429. https://doi.org/10.24963/ijcai.2019/903.
7. Lavanchy, M.; Reichert, P.; Narayanan, J.; Savani, K. Applicants' Fairness Perceptions of Algorithm-Driven Hiring Procedures. *J Bus Ethics* **2023**, p. 125–150.
8. Paramita, D.; Okwir, S.; Nuur, C. Artificial intelligence in talent acquisition: exploring organisational and operational dimensions. *International Journal of Organizational Analysis* **2024**, *32*, 108 – 131. https://doi.org/10.1108/IJOA-09-2023-3992.
9. 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 and Society* **2023**. https://doi.org/10.1007/s00146-023-01783-1.
10. Fabris, A.; Baranowska, N.; Dennis, M.J.; Graus, D.; Hacker, P.; Saldivar, J.; Borgesius, F.Z.; Biega, A.J. Fairness and Bias in Algorithmic Hiring: a Multidisciplinary Survey. *ACM Transactions on Intelligent Systems and Technology* **2024**. https://doi.org/10.1145/3696457.
11. Prasad, K.D.V.; De, T. Generative AI as a catalyst for HRM practices: mediating effects of trust. *Humanities and Social Sciences Communications* **2024**, *11*, 1362. https://doi.org/10.1057/s41599-024-03842-4.
12. Tsiskaridze, R.; Reinhold, K.; Jarvis, M. INNOVATING HRM RECRUITMENT: A COMPREHENSIVE REVIEW OF AI DEPLOYMENT. *MARKETING AND MANAGEMENT OF INNOVATIONS* **2023**, *14*, 239–254. https://doi.org/10.21272/mmi.2023.4-18.
13. Aleisa, M.; Alshahrani, M.; Beloff, N.; White, M. TAIRA-BSC - Trusting AI in Recruitment Applications through Blockchain Smart Contracts. In Proceedings of the 2022 IEEE INTERNATIONAL CONFERENCE ON

 doi:10.20944/preprints202603.1274.v1

BLOCKCHAIN (BLOCKCHAIN 2022), 2022, pp. 376–383. 5th IEEE International Conference on Blockchain (Blockchain), Espoo, FINLAND, AUG 22-25, 2022, https://doi.org/10.1109/Blockchain55522.2022.00059.

14. Kinger, S.; Kinger, D.; Thakkar, S.; Bhake, D. Towards smarter hiring: resume parsing and ranking with YOLOv5 and DistilBERT. *Multimedia Tools and Applications* **2024**, *83*, 82069 – 82087. https://doi.org/10.1007/s11042-024-18778-9.
15. 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 Economic* **2024**, *26*, 421–439. https://doi.org/10.24818/EA/2024/66/421.
16. Schloetzer, J.D.; Yoshinaga, K. Algorithmic Hiring Systems: Implications and Recommendations for Organisations and Policymakers. In Proceedings of the Gill-Pedro, E., Moberg, A. (eds) YSEC Yearbook of Socio-Economic Constitutions 2023: Law and the Governance of Artificial Intelligence, Cham, 2024; pp. 213–246. https://doi.org/10.1007/16495_2023_61.
17. 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:cs.SE/2006.02763].
18. 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, 2020, pp. 1–6. https://doi.org/10.23919/CISTI49556.2020.9140972.
19. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design Science in Information Systems Research. *MIS Quarterly* **2004**, *28*, 75–105. https://doi.org/10.2307/25148625.
20. Mahjoub, A.; Kruyen, P.M. Efficient recruitment with effective job advertisement: an exploratory literature review and research agenda. *International Journal of Organization Theory & Behavior* **2021**, *24*, 107–125. https://doi.org/10.1108/IJOTB-04-2020-0052.
21. 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 preprint arXiv:2302.13971* **2023**.
22. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Łukasz Kaiser.; Polosukhin, I. Attention is all you need. *Advances in neural information processing systems* **2017**, *30*.
23. Boonstra, L. Prompt Engineering. Technical report, Google, 2025.
24. Paradis, T. Entry-Level Jobs in Areas Like Tech Often Require Years of Experience. Business Insider. Url: https://www.businessinsider.com/entry-level-jobs-tech-roles-require-years-of-experience-2024-10.
25. Mohr, T.S. Why Women Don't Apply for Jobs Unless They're 100% Qualified. Harvard Business Review. Url: https://hbr.org/2014/08/why-women-dont-apply-for-jobs-unless-theyre-100-qualified.
26. Johnson, D.; Menezes, A.; Vanstone, S. The Elliptic Curve Digital Signature Algorithm (ECDSA). *International Journal of Information Security* **2001**, *1*, 36—-63. https://doi.org/10.1007/s102070100002.
27. 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, 2023, Vol. 12636, p. 126365H. https://doi.org/10.1117/12.2675431.