This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication.
Citation information: DOI TBC

A Sustainable ML-Based Predictive Model for Semantically Intelligent Search Optimisation in Recruitment Systems

Mohammed-Hassan Ajjam, *Member, IEEE*, and Hamed S. Al-Raweshidy, *Senior Member, IEEE*

Abstract—This paper presents a machine learning-based predictive model that integrates semantic technology and artificial intelligence (AI) to optimise real-time job search and job-resume matching. The proposed model leverages natural language processing (NLP) and semantic feature extraction to improve contextual understanding between job descriptions and candidate resumes. By integrating multiple machine learning algorithms, the system supports accurate matching, efficient ranking, and real-time applicability for recruitment platforms. This study conducts an experimental evaluation using 660 data science job positions from Glassdoor and 962 resumes from Kaggle, representing 25 professional domains. Results show that Support Vector Machines (SVM) effectively filter irrelevant resumes, Random Forest achieves the lowest average mean squared error (MSE) across domains, and XGBoost achieves the lowest runtime, achieving a minimum MSE of 8.903×10^{-5} with an execution time of 0.255 seconds. These results highlight trade-offs between accuracy, ranking effectiveness, and computational efficiency. Overall, the proposed predictive model offers a scalable, fair approach to job-resume matching, supporting data-driven and efficient recruitment decision-making.

Index Terms—Machine Learning (ML), Predictive Models, Real-Time Job Search, Natural Language Processing (NLP), Recruitment Optimisation, Semantic Technology, AI.

1 Introduction

MACHINE learning (ML) field has experienced substantial growth in recent years, leading to notable improvements in predictive modelling across multiple industry domains [1]. One area where these advances have demonstrated particular relevance is job search and recruitment systems [2]. As more users rely on web-based platforms to apply for jobs, the need for predictive models that deliver timely, targeted job search results aligned with user preferences and qualifications is growing [3]. At the same time, large-scale recruitment platforms demand considerable computational resources, raising concerns about scalability and long-term suitability [4, 5]. Consequently, developing efficient ML-based solutions that reduce computational overhead and energy consumption has become essential for sustainable digital recruitment infrastructures [6].

The recruitment landscape is highly dynamic, with organisations continually seeking improved methods to optimise hiring workflows and identify suitable candidates. Recent work has explored machine learning-based predictive models, combined with semantic technology, to support real-time job search optimisation [7]. By leveraging historical and contextual data, such models can assist recruiters in making more informed decisions and generating relevant job-candidate matches [8].

- Mohammed-Hassan Ajjam is with the Department of Electronic and Electrical Engineering (EEE), Brunel University of London, Uxbridge UB8 3PH, London, United Kingdom (e-mail: hassan.ajjam@brunel.ac.uk).
- Hamed Al-Raweshidy is with the Department of Electronic and Electrical Engineering (EEE), Brunel University of London, Uxbridge UB8 3PH, London, United Kingdom (e-mail: hamed.al-raweshidy@brunel.ac.uk).

Beyond improving matching accuracy, these approaches support resource-aware, AI-driven decision-making, enabling informed and efficient recruitment practices [4, 5]. This study examines the design and evaluation of a predictive recruitment model and analyses its potential value for organisations operating in competitive hiring environments [9].

Predictive models that integrate machine learning with semantic technology also support personalised job recommendations and data-driven recruitment analytics [10]. These capabilities can reduce time-to-fill roles and enhance candidate engagement by tailoring search results to individual profiles and preferences [11]. Rather than redefining recruitment practices in abstract terms, such systems provide practical mechanisms to improve the candidate experience and support evidence-based talent acquisition strategies [12]. When applied at scale, data-driven analytics can provide actionable insights to inform workforce planning and hiring decisions in a competitive labour market [13].

This paper presents the underlying methodology, experimental evaluation, and resulting implications. In practical terms, organisations may benefit from improved recruitment workflows through predictive job search results and candidate ranking mechanisms informed by semantic and machine learning techniques.

While the individual techniques employed in this study, such as Latent Semantic Analysis (LSA), cosine similarity, and established machine learning models, are well known, the contribution of this work lies in their systematic integration into a unified, real-time recruitment framework, emphasising system-level design and deployment-oriented evaluation. The proposed system evaluates predictive accuracy, ranking effectiveness, runtime efficiency, and computational resource requirements, providing practical insights for model selection and bridging the gap between methodological research and real-world recruitment system requirements, particularly in

Copyright 'For the purpose of open access, the author has applied a 'Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

contexts where resource constraints and operational demands are critical. The paper makes the following main contributions:

- **Machine Learning Enhanced Predictive Modelling:** Improving job-candidate matching accuracy through semantic similarity and predictive analytics
- **Hiring Efficiency and Computational Cost Reduction:** Enhancing recruitment efficiency while reducing unnecessary computation in large-scale systems.
- **Personalised Candidate Experience:** Supporting individualised job recommendations to improve candidate engagement.
- **Strategic Talent Acquisition**: Providing data-driven insights to assist recruitment professionals with workforce planning and retention strategies.

Overall, this paper demonstrates the practical application of machine learning and semantic technologies in recruitment systems, presenting methods, results, and implications that can support more effective and sustainable hiring practices across diverse organisational contexts.

Further, this paper is structured as follows: Section II reviews related work in HR technology, machine learning-based recruitment systems, and predictive hiring models. Section III describes the proposed methodology, including data selection, preprocessing, model training, evaluation settings, and ethical considerations. Section IV presents the experimental results and discusses matching accuracy, real-time recommendations, and system performance. Section V outlines the implications and suggestions for future research. Finally, Section VI concludes the paper and summarises key findings.

2 Literature Review

Machine learning has become a central component of many modern industrial systems, with applications ranging from supply chain automation and fraud detection to automated sorting and classification tasks [2]. Within this broad landscape, job search and recruitment systems have increasingly adopted machine learning techniques to improve the accuracy and efficiency of job-candidate matching [14]. As web-based recruitment platforms continue to expand, predictive models that automate and optimise search results based on user preferences and qualifications have become essential [15].

2.1 Machine Learning in Hiring

Recruitment processes have evolved in response to the growing availability of data and competitive labour markets, prompting organisations to adopt more systematic and data-driven hiring approaches [16]. Machine learning-based predictive models have been explored to generate real-time recruitment search results by learning patterns from historical hiring data [17]. These models support decision-making by identifying trends and relationships within large datasets, enabling recruiters to assess candidate suitability more consistently and at scale [18].

2.2 Predictive Modelling Key Stages and Advantages

Predictive modelling in recruitment improves talent acquisition by supporting data-driven evaluation and personalisation [19]. Combining predictive analytics with personalised recommendations enables more robust decision-making for recruiters, improves time-to-fill roles and increases candidate satisfaction [20]. Prior research has also demonstrated that reliance on heuristic decision-making has diminished as machine learning models have provided more accurate and cost-effective candidate-matching strategies [21]. These findings highlight the potential of predictive modelling to enhance both operational efficiency and recruitment outcomes.

2.3 Improving the Candidate Experience

Machine learning techniques enable personalised job recommendations by modelling candidate preferences, qualifications, and behavioural data [22]. Such personalisation can improve user engagement and satisfaction by presenting candidates with more relevant opportunities. Empirical studies indicate that tailored recommendation strategies lead to more effective job-seeking experiences compared to generic search mechanisms [22].

2.4 Strategic Talent Acquisition

Beyond operational improvements, machine learning enables HR to leverage real-time data insights to inform strategic talent acquisition decisions. Data analytics represents a data renaissance, providing tools that shape workforce planning and retention. Research has shown that predictive analytics in recruitment helps organisations anticipate hiring needs and respond quickly to market shifts [23].

2.5 Heuristic Techniques vs Predictive Modelling

Heuristic techniques in recruitment often rely on rule-based approaches derived from experience and intuition [24]. While these techniques can be effective, they may lack the scalability and precision offered by predictive modelling. Predictive modelling uses machine learning algorithms to identify patterns in historical data and make data-driven predictions about future hiring needs and candidate fit [25]. This results in accurate and efficient recruitment processes [25]. Unlike static heuristic methods, predictive models can continuously learn and adapt to new data, improving performance over time [15].

2.6 Semantic Technology and AI Integration

Semantic technology, in combination with AI-based recruitment systems, enhances the interpretation of job descriptions and candidate profiles by incorporating contextual and relational information [26]. These approaches enable algorithms to move beyond keyword matching and capture deeper semantic relationships within recruitment data [27]. As a result, semantic integration supports more accurate matching, optimises search results, and improves understanding of job-candidate alignment across diverse hiring scenarios [28].

2.7 Practical Applications and Future Research

This paper extends current research by exploring how to apply machine learning, semantic technology, and AI in recruitment [28]. It details the approach, results, and lessons learned from using machine learning to build a real-time job-search prediction model [29]. These findings are significant because they have several applications across various areas and serve as a basis for future research and practices in machine learning recruitment [29].

In summary, integrating machine learning into recruitment enhances predictive analytics models, improves hiring efficiency, enables greater personalisation of candidate experiences, and supports a more competitive, agile talent acquisition strategy [23]. This paper aims to provide a concrete illustration of real-time applications, practical benefits, and measurable results from machine learning in the context of recruitment, presenting organisations with data-driven conclusions they can leverage to optimise their hiring process within today's rapidly evolving environment [29]. Further studies can examine AI's concerns and suitability in hiring to ensure fair access to employers [30].

2.8 Comparison with Existing Work

What makes our work different from previous work is the following aspects:

1) **Semantic Technology Integration:** Most systems rely primarily on machine learning techniques; the proposed framework explicitly integrates latent semantic feature extraction to improve contextual alignment between job descriptions and resumes. LSA-based representations capture term relationships beyond surface-level keyword matching, enabling more structured semantic alignment during matching.
2) **Real-Time Optimisation:** Existing approaches often depend on offline preprocessing and periodic database updates, limiting responsiveness to new resumes or job postings. In contrast, the proposed framework supports concurrent updating and retrieval of candidate and job data, enabling real-time optimisation.
3) **Candidate Experience Personalisation:** Previous work frequently treats personalisation as a post-processing step. The proposed system integrates semantic similarity and predictive modelling directly into the ranking process to provide structured, data-driven recommendations.
4) **System-Level Design:** Rather than focusing solely on algorithmic performance, the framework integrates semantic modelling, predictive learning, ranking evaluation, runtime analysis, and ethical considerations within a unified architecture for deployment.
5) **Scalability and Practical Applications:** The proposed framework is designed for deployment on commodity hardware using open-source tools, demonstrating scalability under realistic operational constraints.

By addressing these dimensions, this work extends prior studies and provides a practical, system-oriented response to the evolving needs of contemporary recruitment challenges [31].

3 Methodology

This section describes the process used to develop and evaluate a real-time machine learning-based predictive model for job search optimisation using semantic representations and AI techniques. The proposed workflow produces accurate and computationally efficient suitability predictions between job descriptions and candidate resumes under realistic recruitment conditions.

3.1 Data Collection

This study utilises diverse datasets, including candidate resumes and job descriptions sourced from Glassdoor and Kaggle platforms, which are real-world relevant and accurately reflect the current recruitment market. The Job description data, originally from Glassdoor, is available on Kaggle [32] and includes 660 data science job postings. Similarly, resume data is available on Kaggle [33] and comprises 962 resumes across 25 domains (e.g., data science, arts, civil engineering).

The collected datasets consist of structured and unstructured textual features extracted from job descriptions and candidate resumes. Job description features include job title, required skills, responsibilities, experience level, qualifications, industry domain, and job location. Resume features include candidate skills, educational background, work experience, professional domain, and preferences. After preprocessing and filtering, the study retained 481 job descriptions and 166 resumes for experimentation. Pairwise combinations of job descriptions and resumes resulted in 79,846 instances, which were subsequently divided into training and testing sets using a 70:30 split to ensure balanced domain representation.

TABLE I
DATASET DESCRIPTION AND TRAIN/TEST SPLIT

Dataset Component	Description	Count
Job Descriptions	Data science-related job postings (titles, skills, responsibilities, qualifications, domain, location)	481
Resumes	Candidate resumes across 25 professional domains (skills, education, experience, expertise)	166
Total Pairwise Instances	Job-resume combinations used for model training and evaluation	79,846
Training Set (70%)	Instances used for model learning	55,892
Testing Set (30%)	Instances used for model evaluation	23,954

Summary of dataset components, extracted feature types, total sample size, and train-test partitioning used for model training and evaluation.

3.2 Machine Learning Predictive Model Framework

This subsection outlines the internal modelling pipeline, encompassing preprocessing, semantic feature extraction, dataset construction, and predictive learning used to compute job-resume suitability scores.

1) **Data Preprocessing:** Text cleaning is applied to standardise input documents, including duplicate removal, tokenisation, stop-word removal, stemming, and lemmatisation, to convert the text into a numerical form for use with machine learning algorithms. These steps reduce noise and improve consistency before feature extraction.

2) **Natural Language Processing (NLP) and Semantic Feature Extraction:** The system generates semantic representations using Latent Semantic Analysis (LSA) with dimensionality reduction via Truncated Singular Value Decomposition (SVD) [34]. LSA models relationships between documents and terms by projecting sparse term-document matrices into a lower-dimensional space. Truncated SVD efficiently computes the top-k singular components, suitable for sparse text matrices and reducing computational cost relative to complete decomposition.
   - **Dataset Construction:** After preprocessing, the retained 481 job descriptions and 166 resumes were used to create a paired dataset X, which is a pair of job descriptions and resumes. Each job vector is paired with each resume vector, yielding 79,846 instances (481 x 166). The system computes cosine similarity between the LSA representations of each job-resume pair and uses the resulting similarity score as the target variable Y. However, a validated ground-truth label is not available in the dataset. This limitation is common in recruitment research and has been addressed in prior work [35] by evaluating retrieval and ranking behaviour using resume domain classifications rather than recruiter-verified outcomes. Accordingly, the cosine similarity score used as the target variable in this study should be interpreted in the context of semantic job-resume alignment rather than as a validated hiring outcome.
3) **Machine Learning Predictive Modelling:** The study evaluates three machine learning models to predict suitability scores.
   - **Support Vector Machines (SVM):** Used for classification tasks to ensure that job descriptions best match resumes.
   - **Random Forests (RF):** An ensemble learning approach that builds multiple decision trees and merges them for a more accurate prediction.
   - **XGBoost:** Gradient-boosted trees optimising for predictive performance and scalability [36].

   The study selects these models for their effectiveness with structured, high-dimensional representations and their robustness on moderate-sized datasets. Together, they offer a balance between predictive performance and computational efficiency, making them suitable for real-time deployment. The study implements the models using the Scikit-learn library [37] with standard default configurations to ensure fair comparison and computational consistency. Each model incorporates regularisation through its inherent mechanism: SVM applies margin-based control, Random Forest reduces variance via ensemble averaging, and XGBoost constrains model complexity through shrinkage and tree-depth limits. Generalisation performance is assessed using the predefined 70:30 train–test split and repeated experimental runs with standard error reporting.
3) **Design Rationale and Model Scope:** Neural network and transformer-based models, such as BERT-based representations, were not considered in this study due to their higher computational and energy costs, increased inference latency, and reliance on large-scale labelled datasets for effective task-specific fine-tuning. Moreover, prior studies have raised concerns about bias amplification in transformer-based representations in high-stakes decision contexts [38]. Given the focus of this work on scalable, cost-aware, and fair job-resume matching, lightweight machine learning models, combined with semantic feature extraction, were deemed more appropriate for real-time recruitment environments.
4) **Suitability Score Selection Rule**: In the recruitment machine-learning predictive model, each machine-learning technique yields a suitability score for job description j based on candidate resume r. For each job description $j \in J$ and resume $r \in R$, the predictive model outputs a suitability score SS_{jr}. The system expresses the best match for job description j as:

$$Best\ match_j = \max_{r \in R} SS_{jr} \quad (1)$$

3.3 ML Predictive Model System Architecture

The system architecture of the proposed machine learning-based predictive model is illustrated in Figure 1 and follows a modular and layered design that supports scalability, maintainability, and real-time execution. The architecture consists of five main functional components, each responsible for a distinct stage of the end-to-end recruitment pipeline:

- **Data Ingestion:** Collects job descriptions and candidate resumes from datasets.
- **Preprocessing:** Performs text normalisation and quality assurance to ensure that input data are consistent and suitable for semantic analysis. This module acts as a validation gateway before semantic processing.
- **Semantic Feature Extraction:** Converts preprocessed textual inputs into numerical semantic representations, enabling the assessment of contextual similarity between job descriptions and resumes.
- **Predictive Modelling:** Trains and evaluates SVM, RF, and XGBoost models to learn suitability patterns between jobs and resumes. Evaluation checkpoints are incorporated to ensure model reliability before selection.
- **Decision Support:** Ranks candidates based on predicted suitability scores and produces real-time recommendations for each job description.

TABLE II
ALGORITHM REAL-TIME JOB-RESUME MATCHING

Pseudocode: Algorithm Steps
Input:
J = job descriptions
R = resumes
P = preprocessing rules (dedup, tokenise, stop-words, stem, lemmatise)
k = number of LSA components
M = {SVM, RF, XGBoost}

split ratio p (70/30)
thresholds t_MSE, t_rank

Output:
// The m denotes the selected model*
Ranked resumes for each job *j*, selected model *m**
1) Collect datasets J and R.
2) Preprocess:
Repeat:
- Apply text preprocessing to J and R using P.
- Run a preprocessing quality check (missing fields, empty text, noise).
Until the quality check passes.
3) Semantic feature extraction:
- Build TF-IDF matrix on (J ∪ R).
- Apply Truncated SVD with k components to obtain LSA vectors.
- Extract vector z_j for each job j and z_r for each resume r.
4) Build a paired dataset:
For each job j in J:
For each resume r in R:
x_jr = concat(z_j, z_r)
y_jr = cosine_similarity(z_j, z_r) // Target semantic similarity score
store (x_jr, y_jr)
Split into train/test with a ratio p (domain-balanced).
5) Train and evaluate models:
Repeat:
For each model m in M:
- Train m on the training set
- Predict on the test set
- Compute MSE_m
- Compute ranking metric (rank position distribution / mean rank)
- Measure runtime Time_m
Run evaluation quality check using t_MSE and t_rank
Until the evaluation check passes.
6) Select the best model m using accuracy–ranking–runtime criteria.*
7) Deploy:
For each job j:
For each resume r:
SS_jr = m(concat(z_j, z_r)) // Predicted semantic similarity score*
Rank resumes for job j by SS_jr descending.
8) Monitoring:
- Track performance and drift; retrain or update when needed.

Algorithm 1: Pseudocode of the proposed real-time job-resume matching framework using semantic feature extraction and predictive modelling.

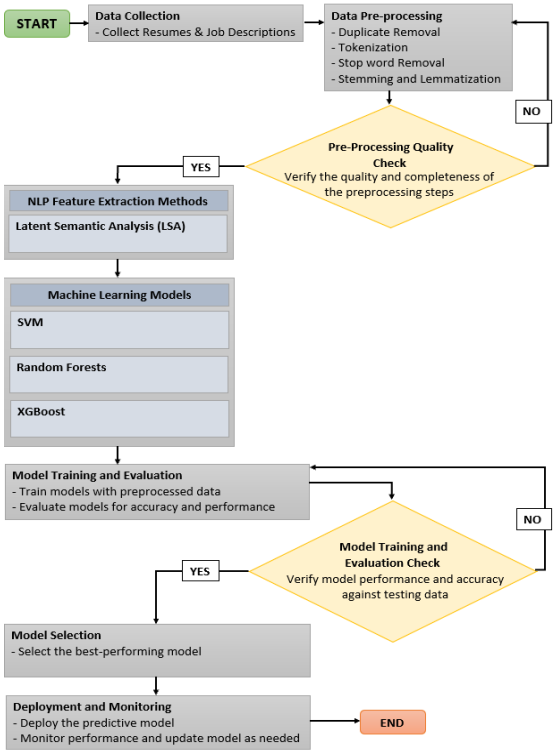

Figure 1: The flowchart of the machine learning predictive model for matching job descriptions with resumes.

3.4 Model Training

The models are trained on the preprocessed paired instances to learn relationships between job descriptions and resumes based on their similarity scores. A train-test split evaluates model performance on test-set job-resume pairings [39].

The curated dataset of 79,846 instances is divided into 70% for training and 30% for testing. To maintain comparable resume-domain representation across both subsets, the split is performed separately within each combination of job ID and resume category. Consequently, individual jobs and resumes may appear in both the training and testing subsets when paired with different counterparts. The evaluation therefore assesses performance on test-set job–resume pairings rather than generalisation to unseen jobs or resumes in the dataset.

The 70:30 split provides sufficient instances for model training while retaining a substantial test subset for evaluation and maintaining consistent resume-domain representation across both subsets. The distribution of resumes across domains is shown in Figure 2, with the Java Developer domain containing the largest number of resumes.

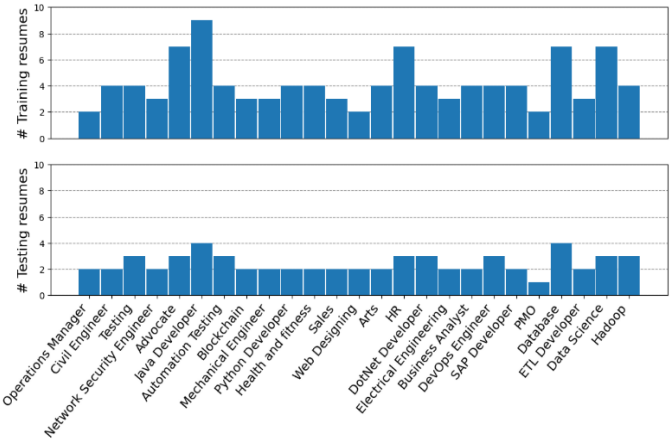

Figure 2: The distribution of resumes across different domains in the training and testing sets.

3.5 Model Evaluation

The experimental evaluation employs quantitative metrics to assess predictive accuracy, ranking effectiveness, and computational efficiency, including Mean Squared Error (MSE), average matching score, resume ranking position, and execution time (runtime). To ensure reliability, *k*-fold stratified random splits were used to evaluate the predictive models, with *k=10*.

The job search optimisation problem is formalised as a regression task by considering the similarity score as the target variable *Y*. Therefore, mean squared error (MSE), as shown in Equation 2, will evaluate the performance of the predictive models, where y_{jr} is the accurate suitability score for job description j by candidate's resumes r, and SS_{jr} is the predictive model output $|R|$ and $|J|$ is the total number of resumes and job descriptions, respectively.

$$MSE = \frac{1}{|R|\cdot|J|}\sum_{j\in J}\sum_{r\in R}(y_{jr} - SS_{jr})^2 \quad (2)$$

Matching job descriptions and resumes based on underlying meaning by employing semantic similarity techniques, such as cosine similarity (Equation 3) and Euclidean distance, measures the similarity between the vector representations of job descriptions and resumes.

Let A be the vector representation of the job description, and let B be the vector representation of the resume. Candidates are ranked based on their similarity scores, where values closer to 1 indicate stronger job-resume matches.

$$Cosine\ Similarity = \frac{\vec{A} \cdot \vec{B}}{\|\vec{A}\| * \|\vec{B}\|} \quad (3)$$

In addition to prediction accuracy and ranking effectiveness, this study measures the execution time of each predictive model to assess suitability for real-time job search optimisation under deployment constraints and resource limitations.

3.6 Model Selection

Comparing the SVM, RF, and XGBoost models and selecting the best-performing one. Let r denote a resume from a relevant domain (e.g., Data Science) for the different Data Science job descriptions. Let r' denote a resume from an irrelevant domain (e.g., HR). The best model is the one that achieves the highest suitability scores for relevant resumes and lower scores for irrelevant resumes, yielding stronger separation in rankings between relevant and irrelevant candidates.

Following C. Yuchen et al. [40], repeat the experiments k times and measure the Standard Error (SE) using Equation 4, where *var* denotes the variance, and MSE_i is the error value of the i^{th} experiment.

$$SE = \frac{\sqrt{var(MSE_1,\ldots,MSE_k)}}{\sqrt{k}} \quad (4)$$

3.7 Ethical and Bias Concerns

The deployment of predictive models in recruitment requires attention to fairness, transparency, and accountability.

1) **Fairness and Bias Mitigation:** This method mitigates bias against gender, ethnicity, and age by excluding sensitive information from the training data, thereby reducing reliance on protected characteristics.
2) **Transparency and Explainability:** The system provides human-interpretable explanations for specific recommendations in plain English and clearly discloses what data it collects, how it uses that data, and the intended purpose of its processing.
3) **Privacy and Data Protection:** The system must protect candidate data from unauthorised access and process it only to the extent necessary, applying data minimisation principles.
4) **Accountability and Governance:** Establish clear accountability for model decisions, including human oversight and the right to challenge algorithmic outcomes. It also involves establishing regular audit points to verify that the model continues to operate ethically and effectively.
5) **Candidate Autonomy and Consent:** Candidates should be informed about automated processing, with options to request a manual review where applicable.

In summary, this section specifies a system-level pipeline for semantic feature extraction, suitability prediction, candidate ranking, and runtime-aware evaluation, with explicit consideration of ethical deployment requirements. Performance outcomes and comparative results are presented in the Results and Discussion section, offering a clearer basis for interpreting the model's overall effectiveness.

4 RESULTS

The experimental evaluation uses four metrics: (i) Mean Squared Error (MSE) to assess prediction accuracy, (ii) average matching score to summarise semantic relevance, (iii) ranking position of the best-matching resume to assess ranking effectiveness, and (iv) execution time (runtime) to assess real-time applicability. All metrics are computed on the test dataset and averaged over repeated runs. MSE values are reported in Table III, matching scores in Table IV, ranking effectiveness in Figure 3, and runtime performance in Figure 4.

4.1 Systematic Evaluation

Conducting a systematic analysis to compare predictive models under controlled settings. For each Data Science job description, the objective is to assign high suitability scores to resumes from relevant domains and low scores to resumes from irrelevant domains. To this end, considering five relevant resume domains, such as Data Science, Python Developer, Database, Hadoop, and ETL Developer, the remaining 20 domains were deemed irrelevant. This grouping reflects the proximity of the domain to Data Science roles and enables a structured assessment of the model's ability to separate relevant from irrelevant candidate profiles.

As shown in Table III, all three predictive models achieved low MSE values and small standard errors (SE), computed across 10 repeated runs ($k=10$) with different random seeds. Random Forest achieved the lowest MSE in most domains (22 domains), while XGBoost achieved the lowest MSE in three domains (Data Science, ETL Developer, and Hadoop). SVM produced higher MSE values than the other two predictive models, but remained within an acceptable range for the evaluated setting.

Table IV reports the average matching score for each resume domain. In relevant domains, higher scores indicate stronger alignment with Data Science job descriptions; in irrelevant domains, lower scores indicate better filtering of non-relevant profiles.

- **XGB** achieved the highest matching score in 2 of 5 relevant domains (Python Developer and Hadoop), with scores of 0.528 and 0.942, respectively. At the same time, it showed the lowest matching scores in 3 out of 20 irrelevant domains (Automation Testing, Electrical Engineering, and Web Designing), with the corresponding values of 0.437, 0.429, and 0.393.

- **Random Forest** achieved the highest matching score in 1 of 5 relevant domains (ETL Developer), with a value of 0.820. However, it performed better than XGB in the irrelevant resume domains by achieving the lowest matching scores in 6 out of 20 irrelevant resume domains (Advocate, Automation Testing, DevOps Engineer, HR, Network Security Engineer, and PMO), with the corresponding values of 0.437, 0.437, 0.433, 0.414, 0.468, and 0.502.
- **SVM** showed a similar performance to XGB. It obtained the highest matching scores in 2 relevant domains (Data Science and Database), with scores of 0.950 and 0.631, respectively. It showed the overall best performance in irrelevant resume domains, obtaining the lowest matching scores in most irrelevant domains (13 domains).

The primary objective of the evaluation based on MSE values is to compare predictive models in terms of their robustness and reliability. In contrast, the assessment based on the matching scores focuses on the performance of the underlying task. All three predictive models demonstrated robustness in achieving the primary objective of the previous evaluation. SVM performed better in the latter review than the other two predictive models, showing stronger separation between relevant and irrelevant domains and providing a more distinct performance advantage overall.

TABLE III
MSE VALUES OBTAINED BY THREE PREDICTIVE MODELS

Resume domain	Predictive models		
	Random Forest	SVM	XGB
Data Science	1.234×10^{-4}	4.351×10^{-3}	7.624×10^{-5}
Python Developer	1.124×10^{-4}	3.124×10^{-3}	1.913×10^{-4}
Database	1.001×10^{-4}	1.018×10^{-3}	1.076×10^{-4}
Hadoop	9.565×10^{-5}	2.786×10^{-3}	7.269×10^{-5}
ETL Developer	1.502×10^{-4}	2.929×10^{-3}	1.297×10^{-4}
Advocate	8.606×10^{-5}	5.082×10^{-3}	1.804×10^{-4}
Arts	8.410×10^{-5}	3.498×10^{-3}	3.337×10^{-4}
Automation Testing	6.350×10^{-5}	1.684×10^{-3}	2.116×10^{-4}
Blockchain	9.017×10^{-5}	1.270×10^{-3}	1.126×10^{-4}
Business Analyst	7.096×10^{-5}	1.639×10^{-3}	1.457×10^{-4}
Civil Engineer	9.014×10^{-5}	2.352×10^{-3}	4.265×10^{-4}
DevOps Engineer	5.737×10^{-5}	2.121×10^{-3}	1.585×10^{-4}
DotNet Developer	9.848×10^{-5}	1.591×10^{-3}	2.009×10^{-4}
Electrical Engineering	7.066×10^{-5}	1.967×10^{-3}	1.956×10^{-4}
HR	6.165×10^{-5}	3.804×10^{-3}	1.624×10^{-4}
Health and Fitness	1.190×10^{-4}	8.944×10^{-3}	2.441×10^{-4}
Java Developer	7.625×10^{-5}	1.813×10^{-3}	1.749×10^{-4}
Mechanical Engineer	6.535×10^{-5}	1.761×10^{-3}	2.924×10^{-4}
Network Security Engineer	7.764×10^{-5}	1.650×10^{-3}	3.117×10^{-4}
Operations Manager	1.031×10^{-4}	3.552×10^{-3}	6.854×10^{-4}
PMO	6.580×10^{-5}	1.382×10^{-3}	2.380×10^{-4}
SAP Developer	1.531×10^{-4}	1.545×10^{-3}	1.910×10^{-4}
Sales	6.546×10^{-5}	1.643×10^{-3}	1.952×10^{-4}
Testing	9.185×10^{-5}	3.449×10^{-3}	3.341×10^{-4}
Web Designing	6.048×10^{-5}	1.444×10^{-3}	1.530×10^{-4}
All domains	$\mathbf{2.094 \times 10^{-4} \pm 1.714 \times 10^{-6}}$	$\mathbf{2.655 \times 10^{-3} \pm 1.052 \times 10^{-5}}$	$\mathbf{8.903 \times 10^{-5} \pm 1.884 \times 10^{-6}}$

Mean Squared Error (MSE) values obtained by the three predictive models across resume domains. The study reports results as MSE ± SE, where SE represents the standard error calculated across repeated experimental runs.

TABLE IV
THE AVERAGE MATCHING SCORE OBTAINED BY THREE PREDICTIVE MODELS

Relevance	Resume domain	Predictive models		
		Random Forest	SVM	XGB
Relevant	Data Science	0.942	**0.950↑**	0.942
	Python Developer	0.527	0.501	**0.528↑**
	Database	0.619	**0.631↑**	0.621
	Hadoop	0.941	0.918	**0.942↑**
	ETL Developer	**0.820↑**	0.783	0.819
Irrelevant	Advocate	**0.437↓**	0.457	0.438
	Arts	0.516	**0.514↓**	0.520
	Automation Testing	**0.437↓**	0.451	**0.437↓**
	Blockchain	0.459	**0.449↓**	0.458
	Business Analyst	0.481	**0.457↓**	0.477
	Civil Engineer	0.356	**0.342↓**	0.352
	DevOps Engineer	**0.433↓**	0.446	0.436
	DotNet Developer	0.467	**0.455↓**	0.465
	Electrical Engineering	0.430	0.433	**0.429↓**
	HR	**0.414↓**	0.445	0.415
	Health and Fitness	0.578	**0.530↓**	0.576
	Java Developer	0.468	**0.459↓**	0.467
	Mechanical Engineer	0.471	**0.468↓**	0.475
	Network Security Engineer	**0.468↓**	0.471	0.471
	Operations Manager	0.387	**0.376↓**	0.382
	PMO	**0.502↓**	**0.502↓**	0.506
	SAP Developer	0.660	**0.637↓**	0.660
	Sales	0.485	**0.461↓**	0.482
	Testing	0.326	**0.289↓**	0.319
	Web Designing	0.397	0.406	**0.393↓**

Average matching scores obtained by the three predictive models across resume domains. Higher values indicate stronger relevance to domain-relevant resumes, while lower values indicate effective filtering of irrelevant resumes.
The best-matching scores for the resume domains are in **bold**.
↑ Denotes the highest score of the relevant resume domain.
↓ Denotes the lowest score of the irrelevant resume domain.

Statistical testing: this study assesses statistical robustness using the Friedman test [41, 42], followed by Holm's post hoc correction [43]. In terms of MSE, Random Forest is significantly better than both XGB and SVM, as evidenced by a p-value of 1.946×10^{-10} at a 5% significance level and evidenced by Holm's post-hoc results in Table V. For average matching scores, all three models performed similarly across relevant and irrelevant domains, with no statistically significant difference, as indicated by p-values of 0.692 and 0.215 at the 5% significance level, respectively. The Friedman test is more appropriate than the ANOVA test, as it is a non-parametric test that does not require normality assumptions, especially when the nature of the problem does not guarantee that similarity scores follow a normal distribution.

TABLE V
HOLM'S POST-HOC CORRECTIONS FOR MSE

Method	avg ranking	z-score	p-value	adj. alpha
Random Forest (Ctrl.)	1.120			
XGB	1.880	2.687	7.210×10^{-3}	0.050
SVM	3.000	6.647	2.991×10^{-11}	0.025

Holm's post-hoc correction results for Mean Squared Error (MSE) across models using Random Forest as the control method.

4.2 Ranking Analysis on Models

Conducting a ranking analysis to evaluate the three predictive models and identifying the best-matching resume for each job description as the one with the highest matching score in the test set. Ranking the resumes based on their

predicted suitability score and recording the rank of the predefined best-matching resume for each predictive model.

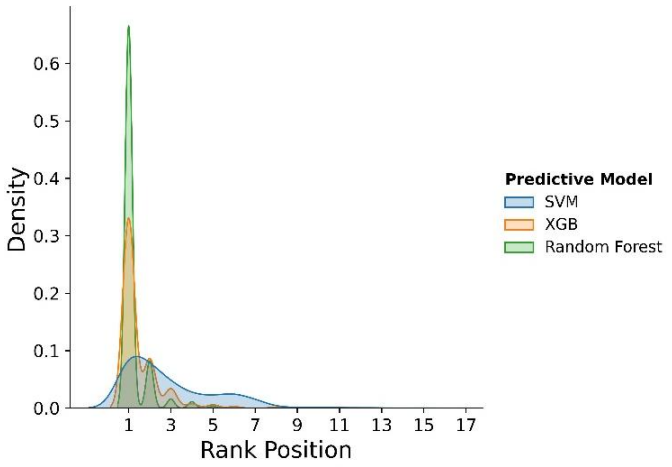

Figure 3: Distribution of ranking positions of the best-matching resume for each job description across the three predictive models.

As illustrated in Figure 3, all three ranking distributions are right-skewed, indicating that the best-matching resume frequently appears in top ranks. Most rankings fall between 1st and 8th position. Random Forest demonstrated the best performance, with an average ranking of 1.208, ranking the best-matching resume between 1st and 5th in most cases. XGB closely followed, with an average ranking of 1.489, ranking the best-matching resume between 1st and 8th. In contrast, SVM had a wider spread, with an average ranking of 2.913, and ranked the best-matching resume between 1st and 5th.

4.3 Model Efficiency and Real-Time Applicability

The predictive models are evaluated in terms of efficiency. Since the aim is to conduct the job-search task in real time, quantifying efficiency by the experiments' runtime is more important than other aspects (e.g., size) in this context. The runtime for the three predictive models was measured by repeating the experiments ten times. Using a machine with an 11th Gen Intel Core i7 processor (CPU) and 16 GB RAM.

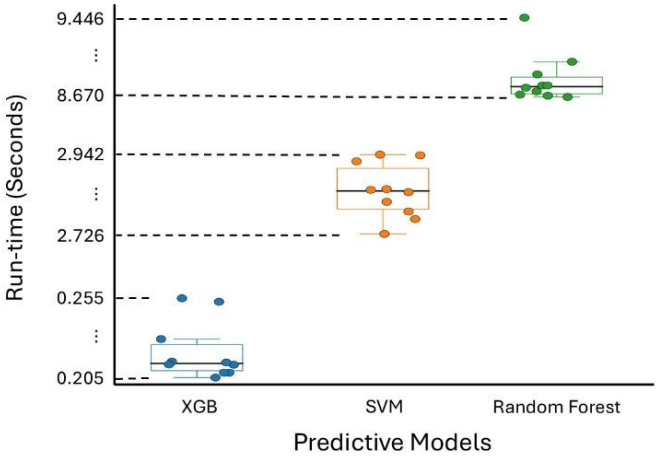

Figure 4: Runtime comparison across predictive models' efficiency measured over ten repeated experimental runs.

Figure 4 shows a substantial runtime difference between Random Forest and the other two models. XGBoost achieved the lowest runtime of 0.255 seconds, followed by SVM, with a median runtime of 2.942 seconds. Random Forest had the highest runtime of 9.446 seconds, indicating lower efficiency and limited applicability for real-time constraints, despite its strong MSE and ranking performance.

TABLE VI
SUMMARY OF EVALUATION METRICS

Metric	Description	Reported In
Mean Squared Error (MSE)	Prediction accuracy of suitability scores	Table III
Average Matching Score	Semantic relevance between jobs and resumes	Table IV
Ranking Position	Rank of the best-matching resume	Figure 3
Runtime (seconds)	Time per model execution	Figure 4

The table lists each evaluation metric, its definition, and the section (table or figure) where the corresponding results appear.

Summary of advantages: the results indicate that the three models offer different trade-offs.

- **Random Forest:** Suggests the best in accuracy, achieving the lowest average MSE across all domains, and excels in ranking resumes effectively, but with a higher runtime.
- **SVM:** The best in relevant match scores and filtering out irrelevant resumes, but with higher MSE and less consistent ranking than RF/XGBoost.
- **XGBoost:** Demonstrates the highest runtime efficiency, reaching the lowest MSE for all domains (8.903×10^{-5}) and the fastest execution time of 0.255 seconds, which supports real-time deployment requirements.

Accordingly, model choice depends on operational constraints, such as whether the deployment prioritises strict latency, numerical stability, or separation performance across irrelevant domains in real-time environments with limited computational resources, scalability considerations, and varying workload conditions.

4.4 Ablation Study Analysis

This study conducts an ablation analysis to evaluate the contribution of selected pipeline components by systematically removing or modifying specific stages and observing the resulting performance changes. The analysis focuses on (i) semantic feature extraction and (ii) model selection strategy. Ablation Settings:

- **The Proposed Model:** Preprocessing, Latent Semantic Analysis (LSA), and machine learning prediction.
- **Single-Model Configuration:** This study evaluates each predictive model (SVM, Random Forest, and XGBoost) independently using the full feature set.
- **Without Model Selection Stage:** A fixed model is used across domains, rather than selecting the best-performing model.

The ablation outcomes indicate that omitting the model selection stage reduces performance in certain domains (Table IV), suggesting that adaptive selection is beneficial when performance varies across domain subsets. Consistent with

Section IV-C, XGBoost provided a strong runtime–accuracy balance, whereas Random Forest achieved stronger accuracy at a higher computational cost.

5 DISCUSSION

This study investigates a real-time recruitment framework that combines semantic feature extraction with machine learning models to support job-resume matching in a scalable, real-time deployment environment. The findings indicate that integrating semantic representations with predictive modelling improves the alignment between job requirements and candidate profiles while maintaining suitability for real-time deployment. Prior research has highlighted the potential of AI-driven recruitment systems to enhance matching accuracy [44]. This work extends prior studies by incorporating a system-level analysis of trade-offs among accuracy, ranking quality, and computational efficiency.

The evaluations provide a systematic analysis of the predictive models across multiple performance dimensions. Tables III and IV, along with Figure 3, show the performance of each model and demonstrate their efficiency for job-matching tasks. Figure 4 evaluates the runtime efficiency, indicating that the proposed model supports scalable, real-time applications. The evaluation results show unique strengths and weaknesses of each predictive model across several key metrics, indicating which model best addresses real-time job matching while balancing accuracy, efficiency, and scalability. The ablation study further demonstrates the contribution of semantic feature extraction and adaptive model selection, as removing these components results in measurable performance degradation.

5.1 Semantic Representation and Model Design Trade-Offs

The Suitability Score (SS) is computed in this study using cosine similarity over latent semantic representations driven via Latent Semantic Analysis (LSA), providing an efficient and interpretable measure of relevance that captures contextual alignment beyond surface-level keyword matching. However, it may not fully model deeper semantic relationships or implicit contextual dependencies present in complex textual data.

Recent transformer-based embedding models, such as SBERT and BERT-based sentence representations, have demonstrated stronger performance in capturing contextual semantics. However, the objective of this study is not to benchmark embedding architectures, but to evaluate a deployment-oriented recruitment framework under real-time and computational constraints. Transformer-based models typically incur substantially higher computational and energy costs, increased inference latency, and greater sensitivity to dataset scale and bias amplification, which may limit their suitability for scalable, resource-aware recruitment systems. Accordingly, this study prioritises lightweight semantic representations that support efficient real-time deployment, particularly in resource-constrained, time-sensitive environments. The integration and benchmarking of contextual embedding models within the proposed framework remain important directions for future work.

5.2 Real-Time Deployment and Computational Efficiency

Beyond predictive performance, the runtime analysis provides insight into the deployment suitability of recruitment models in real-time environments. Models with lower inference latency, particularly XGBoost and SVM, are more suitable for systems that require immediate responses to user interactions. While Random Forest demonstrates strong accuracy and ranking effectiveness, its higher runtime may restrict its use in latency-sensitive applications.

From a sustainability perspective, computational efficiency is often associated with reduced energy consumption in software-based systems. Shorter execution times may correspond to reduced CPU utilisation and lower energy consumption, particularly on large-scale recruitment platforms. This study does not include direct measurements of energy or carbon footprint, and the relationship between runtime and energy use is not empirically verified. The runtime results suggest that lightweight models may contribute to more computationally efficient and potentially more sustainable digital recruitment infrastructures, consistent with prior work on sustainable computing practices [4, 5].

This study addresses the key contributions outlined in the introduction, as follows:

1. **Machine Learning Enhanced Predictive Modelling:** The results show that machine learning improves the accuracy of job candidate selection. Random Forest achieves the lowest average MSE across all domains, showing the highest overall accuracy. XGB reports the lowest MSE values in specific domains, such as Hadoop (7.269×10^{-5}), indicating high predictive accuracy.
2. **Improving Hiring Efficiency:** The implication of real-time optimisations ensures better hiring efficiency. XGB's execution speed (0.255 seconds) makes it a suitable choice for real-time applications, reducing fill time and improving resource efficiency in a real-time digital environment. Although Random Forest demonstrated effective ranking performance, its longer runtime (9.446 seconds) makes it unsuitable for applications that require real-time responses.
3. **Candidate Experience, Personalised:** Using semantic technology, the model fine-tunes the candidate's experience by providing precise, pertinent job suggestions. SVMs produced higher matching scores in relevant domains such as Data Science (0.950) and Database (0.631). Additionally, among the 13 unrelated domains, SVM consistently produced the lowest scores, indicating its effectiveness in filtering out irrelevant resumes and helping candidates receive more targeted job recommendations.
4. **Strategic Talent Acquisition:** These models provide data-driven insights that allow stakeholders to develop effective workforce strategies. Recognising Random Forest's overall accuracy, ranking effectiveness, XGB's runtime efficiency, and domain-specific accuracy enhances recruitment and makes hiring more strategic.

Rather than prescribing a single optimal model, the results suggest that recruitment systems may benefit from adaptive or hybrid strategies that balance accuracy, ranking quality, and runtime according to specific operational requirements and deployment contexts.

5.3 Limitations

First, as discussed in Section III.B.1.b, this study does not use recruiter-confirmed hiring outcomes or independently validated job-resume suitability labels. Instead, it uses resume domain classification as a relevance criterion for evaluating retrieval and ranking behaviour. Consequently, readers should interpret the framework's suitability scores as indicators of semantic job-resume alignment rather than confirmed hiring suitability.

Second, although the datasets are derived from real-world sources (Glassdoor and Kaggle) and reflect authentic job descriptions and candidate resumes, the experimental evaluation relies on a curated subset that may not fully capture the scale, heterogeneity, and continuous data stream of large commercial recruitment platforms.

Third, the study focuses primarily on data science roles, limiting the generalisability of its findings to other professional domains. Fourth, while semantic feature extraction using LSA-based methods improves contextual alignment, it may not fully capture deeper linguistic nuances or implicit competencies embedded in complex recruitment texts.

Additionally, the evaluation does not incorporate dynamic user feedback, recruiter interaction, or longitudinal hiring outcomes, which could further enhance personalisation and adaptive learning in real-world deployment. Finally, although this study explicitly measured runtime efficiency to assess real-time applicability, it did not measure energy consumption directly at the hardware level, such as power usage or energy draw. As a result, sustainability implications are discussed in relation to computational efficiency rather than through explicit energy profiling. Direct carbon-footprint estimation and energy-based benchmarking against transformer-based models were outside the scope of this study, which centres on deployment-oriented predictive modelling rather than energy measurement; these are identified as directions for future work.

5.4 Hardware and Software Cost Analysis

This study presents a real-time machine-learning-based predictive model for cost-efficient deployment on commodity hardware. All experiments were conducted on an 11th-generation Intel Core i7 processor with 16 GB RAM, demonstrating that the proposed framework does not require specialised or high-performance computing infrastructure, such as GPUs or distributed clusters. This configuration reflects realistic deployment conditions for small-scale to medium-scale recruitment platforms.

From a software perspective, the framework relies exclusively on open-source libraries for preprocessing, semantic modelling, and machine learning. Using established techniques such as LSA, SVM, Random Forest, and XGBoost eliminates licensing costs and simplifies system integration. Moreover, adopting computationally efficient models, particularly XGBoost, reduces operational costs by minimising inference time during real-time system operation and large-scale deployment scenarios. Together, these characteristics support deploying scalable, cost-efficient recruitment systems across diverse organisational environments and platforms.

5.5 Comparison with Related Studies

This study compares representative recruitment and job-matching approaches discussed in the Related Work section. As existing studies employ different datasets and evaluation protocols, the comparison focuses on methodological scope and system capabilities rather than direct numerical performance benchmarking across heterogeneous experimental settings. This structured comparison enables a clearer assessment of architectural design choices, the depth of evaluation, and deployment considerations found across prior work.

TABLE VII
COMPARISON WITH RELATED RECRUITMENT STUDIES

Study	Application Focus	Dataset Type	Core Method	Semantic Modelling	Real-Time Evaluation
[7] Al-Quhfa et al.	Talent recruitment	Recruitment data	ML comparison	No	No
[14] Oihab et al.	Intelligent recruitment	Qualitative case-study data	Qualitative case analysis	No	No
[22] Rajput et al.	Resume–job matching	Resume dataset (experimental)	ML recommendation	Partial	No
[23] Van Esch et al.	AI recruitment perceptions	Survey-based experimental data	AI perception modelling (survey-based)	No	No
[43] Sebastien et al.	Fair recruitment matching	Recruitment datasets	Ethical AI matching	Partial	No
Proposed Work	**Real-time job–resume matching**	**Kaggle + Glassdoor**	**Semantic + ML**	**Yes (LSA)**	**Yes**

The table compares representative recruitment and job-resume matching studies with respect to application focus, dataset type, core method, semantic modelling, and real-time evaluation capability. “Partial” denotes the inclusion of contextual or NLP-based components without a fully integrated semantic modelling framework.

Compared to existing studies, the proposed framework integrates semantic feature extraction with multiple machine learning models while explicitly evaluating real-time performance. Although prior work demonstrates the effectiveness of AI and predictive analytics in recruitment, most studies focus on offline accuracy or recommendation relevance, without measuring runtime efficiency or scalability considerations. In contrast, this study evaluates prediction accuracy, ranking effectiveness, execution time, and component contributions through ablation analysis, making it better suited to real-time, resource-aware recruitment systems operating in practical, computationally constrained environments.

6 CONCLUSION

This paper presents a real-time recruitment framework that integrates semantic feature extraction with machine-learning–based predictive modelling to support job-resume matching under practical deployment constraints. By combining Latent Semantic Analysis (LSA) with established machine learning models, the proposed system addresses both matching effectiveness and computational efficiency in dynamic recruitment environments while ensuring more consistent performance across varied application scenarios.

The experimental evaluation demonstrated that different predictive models offer complementary strengths. Support Vector Machines (SVM) showed stronger separation between relevant and irrelevant candidate profiles. Random Forest

achieved robust numerical accuracy and consistent ranking performance. XGBoost achieved the lowest MSE and the fastest runtime, which is well suited to real-time applications. These findings confirm the model's adaptability, allowing stakeholders to select strategies that best suit their needs.

From an operational perspective, the proposed framework supports improved hiring efficiency by reducing unnecessary computation and enabling timely candidate recommendations. The use of lightweight semantic representations and efficient machine learning models makes the system suitable for deployment on commodity hardware, thereby enabling cost-aware, scalable recruitment solutions. Because feature exclusion alone cannot fully ensure fairness, the proposed framework supports responsible deployment through transparency, interpretability, and human oversight.

Future work will extend the proposed predictive model to additional job domains and industry sectors to further evaluate its generalisability. Incorporating dynamic user feedback and recruiter interactions into the learning process represents another promising direction for enhancing personalisation and adaptive decision-making. Moreover, future studies will explore enhanced semantic representations that build on the current framework, with particular emphasis on maintaining transparency, computational efficiency, and real-time applicability. Finally, a more fine-grained sustainability analysis, including direct measurement of energy consumption and deployment in large-scale recruitment environments, will be explored to strengthen the model's applicability in real-world settings.

REFERENCES

[1] T. Hasan and T. Meisen, "Machine learning and deep learning based predictive quality in manufacturing: a systematic review," *Journal of Intelligent Manufacturing*, vol. 33, no. 7, pp. 1879-1905, 2022.

[2] S. Garg, S. Sinha, A.K. Kar, and M. Mani, "A review of machine learning application in human resource management," International *Journal of Productive and Performance Management*, vol. 71, no. 5, pp. 1590–1610, 2022, doi: 10.1108/IJPPM-08-2020-0427.

[3] J. Kenneth et al., "A framework for making predictive models useful in practice," *Journal of the American Medical Informatics Association*, vol. 28, no. 6, pp. 1149-1158, 2021.

[4] A. Mohammad, et al. "Cloud of things (CoT): Cloud-fog-IoT task offloading for sustainable Internet of Things." *IEEE Transactions on Sustainable Computing*, vol. 7, no 1, pp. 87-98, 2020.

[5] Q. Meikang, Sun-Yuan Kung, and Qing Yang. "IEEE transactions on sustainable computing, special issue on smart data and deep learning in sustainable computing." *IEEE Transactions on Sustainable Computing*, vol. 4, no. 1, pp. 1-3, 2019.

[6] IEEE Transactions on Sustainable Computing, "Aims and Scope," *IEEE*, 2023. [Online]. Available: https://www.computer.org/csdl/journal/su/about/15091

[7] Ramachandran, K. K., A. Apsara Saleth Mary, Shibani Hawladar, D. Asokk, Bandi Bhaskar, and J. R. Pitroda. "Machine learning and role of artificial intelligence in optimising work performance and employee behavior." *Materials Today: Proceedings*, vol. 51, pp. 2327-2331, 2022.

[8] H. Al-Quhfa, A. Mothana, A. Aljbri, and J. Song, "Enhancing Talent Recruitment in Business Intelligence Systems: A Comparative Analysis of Machine Learning Models," *Analytics*, vol. 3, no. 3, pp. 297–317, 2024.

[9] S. Agarwal, "Unleashing the Power of Data: Enhancing Physician Outreach through Machine Learning," *International Research Journal of Engineering and Technology*, vol. 10, no. 8, pp. 717-725, 2023.

[10] Karthik, R.V., and Ganapathy S. "A fuzzy recommendation system for predicting customer interests using sentiment analysis and ontology in e-commerce—*Applied Soft Computing*, vol 108, p. 107396, 2021.

[11] C. G. Okatta, F. A. Ajayi, and O. Olawale, "Navigating the future: Integrating AI and machine learning in HR practices for a digital workforce," *Computer Science & IT Research Journal*, vol. 5, no. 4, pp. 1008–1030, 2024.

[12] K. Abdumalik et al., "AI in Human Resource Management: Reimagining Talent Acquisition, Development, and Retention," *International Conference on Knowledge Engineering and Communication Systems (ICKECS)*, vol. 1, pp. 1-8, IEEE, 2024.

[13] S. Guru. "Leveraging AI-Enhanced Analytics for Industry-Specific Optimization: A Strategic Approach to Transforming Data-Driven Decision-Making," *International Journal of Enhanced Research In Science Technology & Engineering*, vol. 10, pp. 78-84, 2021.

[14] AC. Oihab, A. Y. Aranega, and R. C. Sánchez. "Intelligent recruitment: How to identify, select, and retain talents from around the world using artificial intelligence," *Technological Forecasting and Social Change*, vol. 169, p. 120822, 2021.

[15] B. Norhayati et al., "Auto modelling for machine learning: a comparison implementation between rapid miner and python," *Int. J. Emerg. Technol. Adv. Eng.*, vol. 12, no. 5, pp. 15-27, 2022.

[16] Gilch, P.M. and Sieweke, J., "Recruiting digital talent: The strategic role of recruitment in organisations' digital transformation," *German Journal of Human Resource Management*, vol. 35, no. 1, pp. 53-82, 2021.

[17] P. Anand. "November. Machine learning techniques analysis to Efficient resource provisioning for elastic cloud services," *International Conference on Intelligent Computing and Emerging Communication Technologies (ICEC)*, pp. 1-6, IEEE, 2024.

[18] A. B., Deeptanshu, S. Gulati and J. Shukla, "Artificial Intelligence Driven Predictive Analysis of Acoustic and Linguistic Behaviours for ASD Identification," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 11, pp. 5709-5719, 2024, doi: 10.1109/TAI.2024.3439288.

[19] Machireddy, J.R., Rachakatla, S.K. and Ravichandran, P., "Leveraging AI and machine learning for data-driven business strategy: a comprehensive framework for analytics integration," *African Journal of Artificial Intelligence and Sustainable Development*, vol. 1, no. 2, pp. 12-150, 2021.

[20] Okon, R.I.C.H.A.R.D., Odionu, C.S. and Bristol-Alagbariya, B.E.R.N.A.D.E.T.T.E., "Integrating data-driven analytics into human resource management to improve decision-making and organisational effectiveness," *IRE Journals*, vol. 8, no. 6, p. 574, 2024.

[21] D. Bertsimas, and N. Kallus, "The power and limits of predictive approaches to observational-data-driven optimisation," *arXiv preprint arXiv: 1605.02347*, 2016.

[22] A. Rajput, A. Dubey, R. Thakur, D. Singh, et al., "Career craft AI: A Personalised Resume Analysis and Job Recommendations System," *2024 1st International Conference on Innovative Sustainable Technologies for Energy, Mechatronics, and Smart Systems (ISTEMS)*, pp. 1–6, IEEE 2024.

[23] P. Van Esch, J.S. Black, and J. Ferolie, "Marketing AI recruitment: The next phase in job application and selection," *Computers in Human Behavior*, vol. 90, pp. 215–222, 2019.

[24] G. Omid, S. Handley, and R. Stephens. "Logical intuitions or matching heuristic? Examining the effect of deduction training on belief-based reasoning judgments," 2022.

[25] A. Chalfin, O. Danieli, A. Hillis, Z. Jelveh, M. Luca, J. Ludwig, and S. Mullainathan, "Productivity and selection of human capital with machine learning," *American Economic Review*, vol. 106, no. 5, pp. 124–127, 2016.

[26] Q. Lan, D. Wen, Z. Zhang, Q. Zeng, X. Chen, "What is semantic communication? A view on conveying meaning in the era of machine intelligence," *Journal of Communications and Information Networks*, vol. 6, no. 4, pp. 336–371, 2021.

[27] P. Kalyani. "Comprehensive overview of named entity recognition: Models, domain-specific applications and challenges," *arXiv preprint arXiv:2309.14084*, 2023

[28] R. Lobo, P., E. Daga, H. Alani, and M. Fernandez. "Semantic Web technologies and bias in artificial intelligence: A systematic literature review," *Semantic Web*, vol. 14, no. 4, pp. 745-770, 2023.

[29] Y. Jingchen et al.," Energy-aware systems for real-time job scheduling in cloud data centers: A deep reinforcement learning approach," *Computers and Electrical Engineering*, vol 99, p. 107688, 2022.

[30] K. Sara, L. Baker, J. Dolphin, and F. Polli. "AI for hiring in context: a perspective on overcoming the unique challenges of employment research to mitigate disparate impact," *AI and Ethics*, vol. 3, no. 3, pp. 845-868, 2023.

[31] L. Luca, M. Brcic, F. Cabitza, J. Choi, R. Confalonieri, J. Del Ser, R. Guidotti, Y. Hayashi, F. Herrera, A. Holzinger, and R. Jiang. "Explainable Artificial Intelligence (XAI) 2.0: A manifesto of open challenges and interdisciplinary research directions," *Information Fusion*, vol. 106, p. 102301, 2024.
[32] R. Rahman, *Data Science Job Posting on Glassdoor*, Kaggle Dataset, 2020. [Online]. Available: https://www.kaggle.com/datasets/rashikrahmanpritom/data-science-job-posting-on-glassdoor. Accessed: Sep. 1, 2024.
[33] G. Dutta, *Resume Dataset*, Kaggle Dataset, 2019. [Online]. Available: https://www.kaggle.com/datasets/gauravduttakiit/resume-dataset. Accessed: Sep. 1, 2024.
[34] R. Amraoui et al., "NLP and Topic Modeling with LDA, LSA, and NMF for Monitoring Psychosocial Well-being in Monthly Surveys," *Procedia Computer Science*, vol. 251, pp. 398-405, 2024.
[35] MH. Ajjam, and H. S. Al-Raweshidy. "AI-driven semantic similarity-based job matching framework for recruitment systems." *Information Sciences*, 2025: 122728.
[36] R. P. Singh, et al., "December. Comparative Analysis of Tree-Based Models and Deep Learning Architectures for Tabular Data: Performance Disparities and Underlying Factors," *International Conference on Advanced Computing & Communication Technologies (ICACCTech)*, pp. 224-231, IEEE, 2023.
[37] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. and Vanderplas, J. "Scikit-learn: Machine learning in Python." *the Journal of machine Learning research*, vol. *12*, pp.2825-2830, 2011.
[38] Devlin. J, M.W. Chang, K. Lee, and K. Toutanova. "Bert: Pre-training of deep bidirectional transformers for language understanding." In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies,* volume 1 (long and short papers), pp. 4171-4186, 2019.
[39] I. K. Nti, O. Nyarko-Boateng, and J. Aning, "Performance of machine learning algorithms with different K values in K-fold cross-validation," *International Journal of Information Technology and Computer Science*, vol. 13, no. 6, pp. 61–71, 2021, doi: 10.5815/ijitcs.2021.06.05.
[40] C. Yuchen, and Y. Yang. "The one standard error rule for model selection: Does it work?" *Sats*, vol. 4, no. 4, pp. 868-892, 2021, doi:10.3390/stats4040051.
[41] Friedman, M. "The use of ranks to avoid the assumption of normality implicit in the analysis of variance." *Journal of the american statistical association*, vol. 32, no. 200, pp. 675–701, 1937.
[42] Friedman, M. "A comparison of alternative tests of significance for the problem of m rankings." *The annals of mathematical statistics*, vol. 11, no. 11, pp. 86–92, 1940.
[43] Holm, S. "A simple sequentially rejective multiple test procedure." *Scandinavian journal of statistics*, pp. 65–70, 1979.
[44] D. Sebastien et al., "Responsible Artificial Intelligence in Human Resources Technology: An innovative inclusive and fair by design matching algorithm for job recruitment purposes," *Journal of Responsible Technology*, vol. 11, p. 100041, 2022.

Mohammed-Hassan Ajjam (Member, IEEE) received the B.Sc. degree in Computing and the M.Sc. degree in Advanced Computing Technologies from Birkbeck, University of London, London, UK, in 2017 and 2020, respectively, and the PhD degree in Computer Aided Engineering from Brunel University of London, London, UK, in 2026. His research interests include artificial intelligence, semantic technologies, natural language processing, and intelligent information systems.

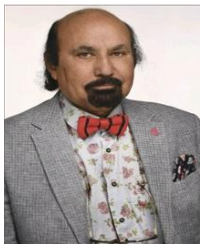

Hamed S. Al-Raweshidy (Senior Member, IEEE) received a Ph.D. from Strathclyde University, Glasgow, U.K., in 1991. Currently, he is a professor in communications engineering. He was with the Space and Astronomy Research Centre, Iraq; PerkinElmer, USA; Carl Zeiss, Germany; British Telecom, U.K.; Oxford U-niversity; Manchester Metropolitan University; and Kent University. He is also the Group Leader of the Wireless Networks and Communications Group (WNCG) and the Director of PG studies (EEE) with Brunel University London, U.K. He is the Co-Director of the Intelligent Digital Economy and Society (IDES), the new research center of the Institute of Digital Future (IDF). He is a course director for the MSc Wireless Communication and Computer Networks. He acts as a consultant and is involved in projects with several companies and operators, such as Vodafone, U.K.; Ericsson, Sweden; Andrew, USA; NEC, Japan; Nokia, Finland; Siemens, Germany; Franc Telecom, Franc; Thales, U.K. and Franc; and Tekmar, Italy, Three, Samsung and Viavi Solutions – actualising several projects and publications with them.