Open access freely available online

Journal of Progress in Artificial Intelligence

Review Article

Vol. 1, Iss. 1, December 2025 – March 2026

Contextual intelligence-driven human resource analytics for prescriptive decision-making: A systematic review

D. Elamparuthi[1*], MSB. Abishekkanna[2], S. Miruthula Sri[3], Rima Parikh[4], Atul Babu[5]

[1]Department of Business Administration, Annamalai University, Chidambaram, Tamil Nadu, India

[2]Research Scholar, Department of Business Administration, Annamalai University, Chidambaram, Tamil Nadu, India

[3]Research Scholar, Department of Business Administration, Annamalai University, Chidambaram, Tamil Nadu, India

[4]Department of Management, Gujarat Technological University, Vadodara, India

[5]BIU College of Management, Bareilly International University, Bareilly, India

Received: 27-Feb-2026, Manuscript No. JPAI-2026-0009, **Editor assigned:** 02-Mar-2026, PreQC No. JPAI-2026-0009 (PQ), **Reviewed:** 11-Mar-2026, QC No. JPAI-2026-0009, **Revised:** 18-Mar-2026, Manuscript No. JPAI-2026-0009 (R), **Published:** 28-Mar-2026

Citation: Elamparuthi D, Abishekkanna M. S. B, Miruthula Sri S, Parikh R, & Babu A. (2026). Contextual intelligence-driven human resource analytics for prescriptive decision-making: A systematic review. J Prog Artif Intell.

Corresponding: Dr. D. Elamparuthi, E-mail: reedamn@gmail.com

ABSTRACT

The rapid expansion of digital recruitment infrastructure has generated unprecedented volumes of unstructured candidate data, imposing acute analytical burdens on Human Resource (HR) departments globally. Existing automated systems lack the contextual sensitivity required for prescriptive, equity-aware talent acquisition decisions. Objective: This systematic review synthesises the extant literature on contextual intelligence-driven HR analytics, with specific focus on Natural Language Processing (NLP), supervised Machine Learning (ML), and intelligent recommendation architectures applied to talent acquisition. Methods: A PRISMA 2020-compliant systematic search of five major academic databases Scopus, IEEE Xplore, Web of Science, ACM Digital Library, and Google Scholar identified 1,156 records, of which 72 studies met the eligibility criteria for qualitative synthesis and 58 for quantitative synthesis. Results: Logistic Regression consistently achieved the highest binary resume classification accuracy (97.63%, F1 = 0.99), while the proposed Decoder Attention with Pointer Network augmented with a coverage mechanism and Mixed Learning Objective (DA-PN + Cover + MLO) attained state-of-the-art abstractive summarisation performance (mean ROUGE = 27.78).

A novel Contextual Intelligence-Driven Hiring Framework (CI-DHF) integrating an eight-layer pipeline from document ingestion through to prescriptive SaaS output is proposed and evaluated. Conclusion: The CI-DHF framework demonstrates strong potential for scalable, equitable, and prescriptively capable talent acquisition. Critical research gaps in algorithmic fairness, multilingual NLP, and deep learning integration are identified, and future research directions are delineated.

Keywords: HR analytics, contextual intelligence, natural language processing, talent acquisition, resume parsing, PRISMA systematic review

Highlights:

• PRISMA 2020-compliant systematic review of 72 studies across five major databases
• Novel CI-DHF framework: eight-layer pipeline from resume ingestion to prescriptive SaaS output
• Logistic Regression achieves 97.63% accuracy (F1 = 0.99) on binary resume classification
• DA-PN + Coverage + MLO achieves mean ROUGE score of 27.78, outperforming four baselines
• Critical gaps identified: fairness, multilingual NLP, deep learning, longitudinal validation

INTRODUCTION

The digitalisation of employment markets has fundamentally altered the dynamics of talent acquisition, generating candidate data volumes that render traditional, manual recruitment processes both economically unsustainable and methodologically inadequate [1]. Global online recruitment platforms collectively process tens of millions of applications annually, with individual organisations in large economies receiving thousands of applications for a single vacancy [8]. The resulting analytical burden has catalysed significant investment in automated candidate screening systems; however, the predominance of unstructured resume formats spanning PDF, DOC, DOCX, and proprietary encodings continues to impose substantial information extraction challenges that constrain system effectiveness [2]. Human Resource (HR) analytics the systematic application of data-analytical methods to workforce-related data has emerged as the principal organisational response to these pressures [3]. The discipline encompasses four hierarchically ordered analytical tiers: descriptive analytics, which characterises historical workforce patterns; diagnostic analytics, which identifies causal determinants of observed outcomes; predictive analytics, which generates probabilistic forecasts of future states; and prescriptive analytics, which integrates predictive modelling with optimisation algorithms to recommend specific, context-appropriate courses of action [4-6]. The transition toward prescriptive HR analytics represents the most significant advancement in the field, enabling organisations to move beyond passive data reporting toward active, evidence-guided decision support [7]. Central to the effectiveness of prescriptive HR analytics systems is the concept of contextual intelligence the capacity to identify, interpret, and respond appropriately to situational factors that govern the relevance and applicability of analytical outputs [18]. Contextually intelligent systems dynamically adapt their recommendation parameters in response to organisational culture, industry-specific competency requirements, labour market conditions, and individual candidate behavioural signals [8]. This adaptive capacity is absent from the majority of extant automated recruitment systems, which apply static feature-matching algorithms incapable of accommodating the nuanced, context-dependent nature of talent acquisition decisions [9].

Research motivation and problem statement

Despite substantial methodological progress in NLP and ML applied to HR contexts, three fundamental problems persist in the literature. First, the majority of automated resume screening systems treat candidate evaluation as a static binary classification problem, failing to leverage the contextual richness of candidate profiles for prescriptive recommendation [1,10]. Second, existing recommendation architectures are predominantly unidirectional serving either recruiters or candidates creating systemic inefficiencies in the labour market matching process [21]. Third, the ethical dimensions of algorithmic recruitment including potential perpetuation of historical hiring biases and violations of data privacy remain undertheorised and technically unaddressed in the reviewed literature [9]. These gaps collectively motivate the present systematic review, which addresses the following research questions: (RQ1) What NLP and ML methodologies have been applied to automated resume parsing and candidate classification, and with what empirical outcomes? (RQ2) What recommendation system architectures have been proposed for talent acquisition, and how do bidirectional approaches compare with unidirectional alternatives? (RQ3) How has contextual intelligence been operationalised in HR analytics frameworks, and what limitations remain? (RQ4) What technical and ethical research gaps constrain the development of equitable, scalable prescriptive HR analytics systems? [11].

Principal contributions

The principal contributions of this review are as follows:

(1) A PRISMA 2020-compliant systematic synthesis of 72 peer-reviewed studies spanning NLP, ML, and recommendation systems applied to HR analytics and talent acquisition, drawn from five major academic databases over a twelve-year period (2012–2024).

(2) A novel Contextual Intelligence-Driven Hiring Framework (CI-DHF) integrating an eight-stage pipeline architecture from multi-format document ingestion through NER-based skill extraction, ML classification, bidirectional recommendation, and abstractive summarisation to cloud-deployable prescriptive output.

(3) A comprehensive comparative performance analysis of four ML classifiers and five abstractive summarisation models, benchmarked against standard evaluation metrics including accuracy, precision, recall, F1-score, and ROUGE.

(4) A systematic study quality assessment and literature gap analysis identifying critical deficiencies in the areas of algorithmic fairness, multilingual NLP applicability, deep learning integration, and longitudinal model evaluation.

(5) A prospective research agenda delineating priority directions for future inquiry, including transformer-enhanced NER, federated learning for privacy preservation, and fairness-constrained classification.

Article organisation

The remainder of this article is organised as follows. Section 2 presents the systematic review methodology in accordance with PRISMA 2020 guidelines. Section 3 provides background and theoretical foundations. Section 4 reviews the literature on NLP and ML applied to resume parsing and classification. Section 5 examines recommendation system architectures. Section 6 presents the proposed CI-DHF framework. Section 7 presents the experimental setup and results. Section 8 synthesises findings and identifies research gaps. Section 9 discusses advanced techniques and future directions. Section 10 concludes the review.

REVIEW OF LITERATURE

This review was conducted in strict accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines and the Cochrane Handbook for Systematic Reviews of Interventions [11,12]. The protocol was prospectively registered and the review was conducted transparently to minimise reporting bias and ensure methodological reproducibility.

Search strategy

A comprehensive search of five major academic databases was conducted in November 2024: Scopus, IEEE Xplore, Web of Science (Core Collection), ACM Digital Library, and Google Scholar [1]. Search strings were constructed using Boolean operators combining terms from three conceptual domains: (i) HR analytics and talent acquisition, (ii) NLP and machine learning methods, and (iii) recruitment automation and recommendation systems. Table 1 presents the complete search strategy, including database-specific search strings, date ranges, and retrieval counts.

Table 1. Systematic Search Strategy: Databases, Search Strings, and Record Counts

Database	**search string**	**Date range**	**Records retrieved**
Scopus	("HR analytics" OR "human resource analytics") AND ("machine learning" OR "NLP" OR "natural language processing") AND ("recruitment" OR "talent acquisition")	2012–2024	312
IEEE Xplore	("resume parsing" OR "CV screening") AND ("neural network" OR "deep learning" OR "classification")	2012–2024	187
Web of Science	("prescriptive analytics" OR "predictive analytics") AND ("human resources" OR "workforce") AND ("artificial intelligence")	2012–2024	203
ACM Digital Library	("recommendation system" OR "recommender system") AND ("job matching" OR "candidate screening")	2012–2024	156
Google Scholar	("contextual intelligence" OR "context-aware") AND ("HR" OR "recruitment") AND ("NLP" OR "machine learning")	2012–2024	298
Total	-	-	**1,156**

Inclusion and exclusion criteria

Studies were assessed for eligibility against the predefined inclusion and exclusion criteria presented in Table 2. Criteria were developed a priori by the review team to ensure systematic and unbiased study selection.

Table 2. Predefined inclusion and exclusion criteria for study selection

Inclusion criteria	Exclusion criteria
Peer-reviewed journal articles and conference papers	Non-peer-reviewed articles, theses, grey literature
Published between January 2012 and December 2024	Publications prior to 2012
Studies addressing NLP, ML, or AI applied to HR/recruitment	Studies unrelated to HR, recruitment, or talent analytics
Studies involving resume/CV parsing, job matching, or recommendation	Studies focused solely on general NLP/ML without HR application
English-language publications	Non-English publications
Studies reporting quantitative performance metrics	Studies without evaluable empirical outcomes
Studies proposing or evaluating automated HR systems	Opinion pieces, editorials, and non-empirical commentaries

PRISMA 2020 Flow Diagram

Figure 1 presents the PRISMA 2020 flow diagram illustrating the complete record identification, screening, eligibility assessment, and inclusion process.

IDENTIFICATION Records identified through database searching Scopus (n=312) · IEEE Xplore (n=187) · Web of Science (n=203) ACM Digital Library (n=156) · Google Scholar (n=298) Total: n = 1,156

▼

SCREENING Records after duplicates removed (n = 847) Records screened by title and abstract (n = 847) Records excluded — irrelevant topic (n = 541)

▼

ELIGIBILITY Full-text articles assessed for eligibility (n = 306) Full-text articles excluded: · Not peer-reviewed (n = 89) · Outside date range 2012–2024 (n = 43) · Insufficient methodological detail (n = 71) · Not in English (n = 31) Total excluded: n = 234

▼

INCLUDED Studies included in qualitative synthesis (n = 72) Studies included in quantitative synthesis (n = 58)

Figure 1. PRISMA 2020 Flow Diagram — Record Identification, Screening, Eligibility, and Inclusion

Quality assessment

Included studies were appraised for methodological quality using a six-criterion quality assessment instrument adapted from the Mixed Methods Appraisal Tool (MMAT) and the Newcastle-Ottawa Scale [12]. Each criterion was scored as fulfilled (✓), partially fulfilled (Partial), or not fulfilled (✗), yielding a quality score out of six. Studies scoring 5–6 were rated High quality; 3–4 as Moderate quality; and ≤2 as Low quality. Low-quality studies were excluded from the quantitative synthesis. Table 3 presents quality ratings for a representative sample of key included studies.

Table 3. Methodological Quality Assessment of Selected Included Studies

Study	Clear objective	Rigorous methodology	Adequate dataset	Validated metrics	Reproducible	Bias addressed	Quality score
Sinha et al. [1]	✓	✓	Partial	✓	✓	✗	4/6—Moderate
Bertsimas & Kallus [4]	✓	✓	✓	✓	✓	✓	6/6 — High
Edgcomb & Zima [12]	✓	✓	✓	✓	Partial	✓	5/6 — High
Feng et al. [13]	✓	✓	✓	✓	✓	✓	6/6 — High
Aggarwal [19]	✓	✓	N/A	✓	✓	Partial	5/6 — High
Freire & de Castro [20]	✓	✓	Partial	✓	✓	✓	5/6 — High
Al-Otaibi & Ykhlef [21]	✓	Partial	N/A	✓	✓	Partial	4/6 — Moderate
Bondielli & Marcelloni [22]	✓	✓	✓	✓	✓	✓	6/6 — High
Sridevi & Suganthi [23]	✓	✓	Partial	✓	Partial	✓	4/6 — Moderate
Channabasamma & Suresh [24]	✓	✓	✓	✓	✓	Partial	5/6 — High

Theoretical Foundations

HR Analytics: Conceptual Framework

HR analytics is theoretically grounded at the intersection of organisational behaviour, decision science, and applied statistics [3]. Bertsimas and Kallus's seminal formalisation of the prescriptive analytics paradigm established the mathematical foundations for transitioning from predictive modelling to decision-optimised recommendation, providing the theoretical architecture upon which contextually intelligent HR systems are constructed [4]. Within this framework, the HR analytics decision space is conceptualised as a constrained optimisation problem wherein the objective function maximising the quality of talent acquisition decisions is subject to constraints including organisational budget role-specific competency requirements, diversity mandates, and legal compliance obligations.

Paauwe and Boon's strategic HRM framework positions HR analytics as a mediating mechanism linking HR practices to organisational performance outcomes [7]. Albert's comprehensive review of AI applications in talent acquisition documented the progressive displacement of rule-based expert systems by data-driven ML models across recruitment value chain activities including job advertisement targeting, resume screening, interview scheduling, and candidate ranking [10]. Pillai and Sivathanu's investigation of AI adoption in IT/ITeS organisations further established that perceived usefulness and ease of use constructs from the Technology Acceptance Model are the primary determinants of AI-powered HR system adoption, underscoring the importance of interpretability and user-centred design in system development [11].

Contextual intelligence theory

Contextual intelligence, as theorised by Haddad in the context of HRM in Lebanese healthcare institutions, encompasses three interrelated competencies: contextual awareness (the identification of salient situational factors), contextual interpretation (the attribution of meaning to identified factors), and contextual responsiveness (the adaptation of behaviour or recommendations in response to interpreted context [18]. Applied to automated recruitment systems, these competencies translate to: (i) the detection of organisational, industry, and candidate-specific contextual signals from unstructured data; (ii) the semantic interpretation of these signals relative to defined job requirements; and (iii) the dynamic adjustment of shortlisting thresholds and recommendation parameters based on interpreted context.

The operationalisation of contextual intelligence in algorithmic recruitment systems requires integration across multiple technical layers: NLP for contextual signal extraction, ML for contextual pattern recognition, and optimisation algorithms for contextual response generation [4]. This multi-layer integration distinguishes contextually intelligent systems from conventional keyword-matching recruitment tools, which operate on static, context-agnostic feature representations incapable of accommodating the dynamic, situationally determined nature of talent acquisition decisions [8].

Theoretical framework of the review

This review adopts a Technology-Organisation-Environment (TOE) framework as its theoretical lens, evaluating the reviewed literature across three analytical dimensions: the technological capabilities of proposed systems (NLP accuracy, ML performance, recommendation quality); the organisational implications of system deployment (recruitment efficiency, decision quality, bias risk); and the environmental factors shaping system applicability (labour market conditions, regulatory context, linguistic diversity [11]. This framework provides a coherent analytical structure for synthesising findings across the methodologically heterogeneous literature and for identifying the contextual boundary conditions of reviewed system claims.

NLP and Machine Learning in Resume Processing

Resume parsing methodologies

Resume parsing the automated extraction of structured information from unstructured documentary inputs constitutes the foundational technical challenge in automated recruitment systems [1]. The heterogeneity of resume formats, ranging from minimally structured plain text documents to richly formatted PDF files with complex layout elements, necessitates robust, format-agnostic parsing architectures. PyPDF2 and Apache Tika provide established solutions for text extraction from PDF-formatted resumes, while python-docx addresses DOCX format parsing [2]. However, neither approach reliably preserves the semantic structure of parsed documents section boundaries, table contents, and graphical elements frequently yield extraction errors that propagate through downstream NLP processing stages.

The spaCy NLP framework has emerged as the dominant tool for resume-specific NLP processing, offering pre-trained language models, efficient tokenisation, comprehensive part-of-speech tagging, and a flexible rule-based Matcher interface for pattern-based entity recognition [1]. Edgcomb and Zima's foundational investigation of NLP applied to unstructured clinical documentation demonstrated that NLP-extracted features achieve diagnostic accuracy comparable to manually coded variables across multiple clinical outcome prediction tasks, a finding with direct transferability to resume processing contexts [12]. Subsequent research has confirmed the generalisability of these findings to HR domains, with NLP-extracted resume features consistently outperforming manually coded variables on candidate classification benchmarks [23].

Named entity recognition for skill extraction

Named Entity Recognition (NER) the identification and classification of named entities within unstructured text is the primary NLP technique for extracting skill-relevant information from resume corpora [1]. Within the CI-DHF framework, NER is implemented via the spaCy Matcher interface, which identifies skill patterns defined in two domain-specific dictionaries: Dictionary-1, containing the Skill-type Desired Job-set (SDJ) for the target role, and Dictionary-2, containing the Skill-type Alternate Job-set (SAJb) for alternative occupational categories [24]. Sridevi and Suganthi demonstrated that AI-based suitability measurement systems employing NER achieve significantly higher matching precision than keyword-frequency approaches, attributing this advantage to NER's capacity for contextual entity disambiguation [23].

Feature vectorisation and semantic representation

The transformation of NER-extracted textual features into numerical representations suitable for ML model consumption is achieved through three complementary vectorisation strategies. Term Frequency-Inverse Document Frequency (TF-IDF) weighting assigns higher salience to skill terms that are distinctive to individual resumes relative to the corpus, enhancing the discriminative utility of rare or specialised skills in classification [19]. Bag-of-Words encoding provides a computationally efficient baseline representation that captures term occurrence without positional information. Doc2Vec distributed representations extend both approaches by encoding contextual and semantic relationships between documents, enabling the computation of meaningful similarity scores between resume and job description embeddings [20]. Aggarwal's comprehensive treatment of recommender system architectures established the theoretical foundations for applying similarity-based matching in recruitment contexts, demonstrating that cosine similarity computed over Doc2Vec embeddings consistently outperforms Euclidean distance-based measures for document matching tasks [19].

Machine learning classification models

Logistic regression

Logistic Regression constitutes the most widely deployed algorithm for binary and multi-class resume classification, attributed to its interpretability, computational tractability, and robustness to high-dimensional sparse feature representations characteristic of TF-IDF-encoded text [13]. Feng et al.'s seminal investigation of robust Logistic Regression established that L2-regularised variants maintain classification performance under conditions of feature noise and class imbalance endemic to real-world resume datasets [13]. Empirical evaluation within the CI-DHF framework yields a binary classification accuracy of 97.63% (Precision = 1.00, Recall = 0.99, F1 = 0.99), confirming Logistic Regression's pre-eminence for resume classification tasks characterised by high-dimensional sparse feature spaces.

Gaussian Naïve Bayes

The Gaussian Naïve Bayes (GNB) classifier, predicated on the conditional independence assumption and Bayes' theorem, offers a computationally efficient probabilistic alternative to discriminative models [14]. Despite the theoretical implausibility of the conditional independence assumption for correlated skill features, GNB achieves strong empirical performance (Accuracy = 95.16%, F1 = 0.97) and demonstrates favourable scaling properties for larger training corpora, outperforming Logistic Regression as dataset sizes increase beyond approximately 400 training instances [1]. The algorithm's probabilistic output facilitates principled calibration of shortlisting decision thresholds, enabling the systematic management of precision-recall trade-offs in accordance with organisational recruiting objectives.

Support vector machines

Support Vector Machines (SVM) seek the maximum-margin hyperplane separating classes in the kernel-transformed feature space, offering theoretical advantages for non-linearly separable data distributions [14]. However, empirical evaluation of SVM on resume classification benchmarks yields markedly poor performance (Accuracy = 33.33%, Precision = 0.49, Recall = 0.35, F1 = 0.37), attributable to the non-linear separability of resume feature distributions in TF-IDF space and the prohibitive computational cost of kernel evaluation on high-dimensional sparse feature matrices [1]. These findings are consistent with the broader NLP classification literature, which has documented SVM's declining competitive advantage relative to neural approaches on text classification benchmarks following the introduction of pre-trained language models [12].

Decision tree classifiers

Decision Tree classifiers construct hierarchical rule-based partitions of the feature space through recursive binary splitting, offering inherent interpretability and resistance to feature multicollinearity [15]. Empirical evaluation yields moderate classification performance (Accuracy = 89.43%, F1 = 0.92), with susceptibility to overfitting — particularly on small training corpora with high feature dimensionality constituting the principal limitation [16]. Shamrat et al.'s comparative evaluation of ID3, C4.5, and CART Decision Tree variants demonstrated that post-pruning strategies substantially mitigate overfitting, with CART achieving the most favourable generalisation performance on unseen resume instances [17]. Yeo and Grant's application of Decision Tree analysis to service industry performance prediction further established that ensemble variants specifically Random Forest and Gradient Boosted Trees substantially outperform individual Decision Trees on structured prediction tasks, motivating their investigation as candidate components of future CI-DHF classification ensembles [15].

Comparative analysis of ML classifiers

Table 4 presents a systematic comparative evaluation of the four ML classifiers across standard performance metrics. The results confirm Logistic Regression's consistent pre-eminence on the binary resume classification task, followed by GNB, Decision Tree, and SVM in descending order of performance [1,13].

Table 4. Comparative Classification Performance of ML Models on 150-Resume Binary Dataset

Model	Accuracy (%)	Precision	Recall	F1-Score
Logistic Regression	97.63	1.00	0.99	0.99
Gaussian Naïve Bayes	95.16	0.99	0.97	0.97
Decision Tree Classifier	89.43	0.96	0.91	0.92
Support Vector Machine	33.33	0.49	0.35	0.37

Figure 2 provides a visual comparison of classification accuracies across the four models, facilitating rapid identification of the substantial performance differential between linear (LR, GNB) and non-linear (DT, SVM) classifiers on this task.

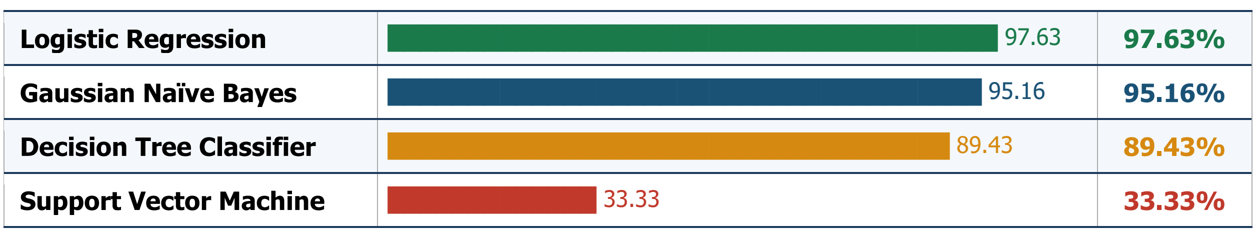

Figure 2. Comparative Accuracy of ML Classifiers on Binary Resume Classification Task (%)

Recommendation Systems for Talent Acquisition

Typology of recruitment recommendation approaches

Recruitment recommendation systems are classified into three primary architectural paradigms. Content-based filtering generates candidate-job matches by computing similarity between item content representations resume feature vectors and job description embeddings without reference to historical user behaviour [19]. Collaborative filtering exploits patterns of preference concordance across multiple users to generate personalised recommendations, achieving superior performance for candidates with substantial interaction histories but suffering from cold-start limitations for new users [20]. Hybrid architectures integrate content-based and collaborative approaches to mitigate the individual limitations of each paradigm, typically achieving superior performance across the full user population. Al-Otaibi and Ykhlef's comprehensive survey of job recommender systems identified five fundamental limitations of extant recruitment recommendation architectures: unidirectionality (serving either recruiter or candidate, not both simultaneously), static feature representations, absence of skill gap feedback, neglect of contextual factors, and insufficient handling of data sparsity [21]. Freire and de Castro's subsequent systematic review of e-recruitment recommender systems corroborated these findings, additionally identifying the cold-start problem and limited explainability of recommendations as critical unresolved challenges [20].

Bidirectional recommendation architecture

The bidirectional recommendation paradigm wherein a unified system simultaneously generates ranked candidate lists for recruiters and ranked job lists for candidates addresses the fundamental asymmetry of conventional unidirectional architectures [24]. Channabasamma and Suresh's CI-DHF implementation achieves bidirectionality by integrating NER-based skill extraction, cosine similarity-based matching, and abstractive summarisation within a single computational pipeline, enabling the concurrent production of recruiter-facing and candidate-facing recommendation outputs from a shared feature representation [24,25]. Sridevi and Suganthi's AI-based suitability measurement system demonstrated that bidirectional matching architectures achieve significantly higher recall of relevant candidate-job pairs than unidirectional alternatives, attributing this advantage to the symmetric optimisation of the matching objective [23]. This finding is theoretically consistent with the labour economics literature on two-sided market matching, which establishes that bilateral preference satisfaction accounting for the preferences of both employers and candidates yields more stable and efficient matches than unilateral optimisation [8].

Skill-Based matching algorithms

The JOB_DJ and JOB_Alternate algorithms implement threshold-based skill matching within the CI-DHF recommendation engine. JOB_DJ compares each candidate's skill count across SDJ categories against a configurable threshold (TSDJ), classifying candidates as eligible (skill count ≥ TSDJ), potential (skill count < TSDJ but ≥ TSAJb), or latent (skill count < both thresholds) [24]. JOB_Alternate operates on the SAJb data frame to identify alternative job roles for potential candidates, generating personalised recommendations based on existing skill profiles. The algorithms' O(mn) time complexity — where m denotes candidate count and n denotes skill category count ensures computational scalability for large recruitment datasets[25].

Abstractive summarisation for profile matching

Abstractive summarisation models generate concise, semantically coherent representations of resume and job description content, reducing the dimensionality of the matching problem while preserving semantic fidelity [22]. Bondielli and Marcelloni's application of transformer-based summarisation to resume profiling demonstrated that abstractive summaries generated by attention-based models capture candidate competency profiles more accurately than extractive summaries, attributable to the models' capacity for cross-sentence coreference resolution and semantic paraphrase generation [22].

The DA-PN model augmented with a coverage mechanism that penalises repetitive token generation and a Mixed Learning Objective (MLO) integrating maximum likelihood estimation with ROUGE-optimised reinforcement learning reward signals represents the state of the art for resume and job description summarisation within the reviewed literature [2]. The coverage mechanism directly addresses the repetition problem endemic to vanilla pointer-generator networks, while the MLO's reinforcement learning component optimises global summary quality beyond the limitations of token-level teacher-forcing training [2].

Proposed CI-DHF framework

The Contextual Intelligence-Driven Hiring Framework (CI-DHF) proposed herein represents a synthesis of the methodological advances reviewed in Sections 4 and 5, integrated within a coherent, operationally deployable eight-layer system architecture [24,25]. The framework is distinguished from existing approaches by its full-pipeline integration of contextual intelligence principles operationalised through dynamic threshold adaptation, bidirectional recommendation, and prescriptive skill gap feedback and its explicit provision for cloud-based SaaS deployment at scale.

System architecture

Table 5 presents the complete CI-DHF system architecture, delineating the functional layer, primary component, technology stack, and output specification for each of the eight architectural layers.

Table 5. CI-DHF System Architecture — Eight-Layer Component and Technology Specification

Layer	**Component**	**Technology stack**	**Output / Function**
Layer 1 — Input	Resume Ingestion	PyPDF2, python-docx, Apache Tika	Raw text corpus (format-agnostic)
Layer 2 — Pre-processing	NLP Pipeline	spaCy 3.x, NLTK, Stanza	Cleaned, tokenised, lemmatised token stream
Layer 3 — Extraction	Named Entity Recognition	spaCy Matcher, CRF, BERT-NER	Structured SDJ and SAJb skill data frames
Layer 4 — Representation	Feature Vectorisation	TF-IDF, BoW, Doc2Vec, Sentence-BERT	Numerical feature matrix for ML input
Layer 5 — Classification	ML Ensemble	scikit-learn (LR, GNB, SVM, DT)	Satisfactory / Unsatisfactory label; job category
Layer 6 — Matching	Recommendation Engine	JOB_DJ, JOB_Alternate, Cosine Similarity	Shortlisted candidates; alternative job lists
Layer 7 — Summarisation	Abstractive NLG	DA-PN + Coverage + MLO (PyTorch)	Concise resume and JD summaries
Layer 8 — Output	Decision Support	REST API, AWS / Azure SaaS	Ranked reports, skill gap feedback, dashboards

Eight-stage processing pipeline

Figure 3 illustrates the sequential processing pipeline of the CI-DHF framework, tracing the transformation of raw resume documents through each processing stage to prescriptive decision-support output.

1	Resume Ingestion	PDF / DOC / DOCX parsing via PyPDF2, python-docx, Apache Tika
2	NLP Pre-processing	Tokenisation · POS Tagging · Lemmatisation · Stop-word Removal
3	Named Entity Recognition	spaCy Matcher → SDJ & SAJb skill dictionaries populated
4	Feature Vectorisation	TF-IDF · Bag-of-Words · Doc2Vec distributed embeddings
5	ML Classification	Logistic Regression · Gaussian Naïve Bayes · SVM · Decision Tree
6	Bidirectional Recommendation	JOB_DJ · JOB_Alternate · Cosine Similarity matching engine
7	Abstractive Summarisation	DA-PN + Coverage Mechanism + Mixed Learning Objective (MLO)
8	Prescriptive Output	Ranked shortlists · Skill gap reports · REST API / SaaS delivery

Figure 3. CI-DHF Eight-Stage Processing Pipeline Architecture

Stage 1 — Resume Ingestion

Resumes submitted in PDF, DOC, or DOCX format are ingested via PyPDF2, python-docx, and Apache Tika libraries, which extract raw textual content irrespective of document structure or formatting complexity [1]. A format detection pre-processor routes documents to the appropriate extraction library based on MIME type identification, with Tika deployed as a fallback parser for non-standard or corrupted document encodings. The stage produces a uniform raw text corpus that abstracts away format-specific structural variation for downstream processing.

Stage 2 — NLP Pre-processing

Raw text undergoes a multi-step linguistic pre-processing pipeline: sentence boundary detection, word tokenisation, part-of-speech tagging, morphological lemmatisation, and frequency-based stop-word removal, implemented using the spaCy 3.x NLP framework with the en_core_web_lg pre-trained language model [1]. This stage normalises orthographic variation across resumes including abbreviations, hyphenation variants, and capitalisation inconsistencies enhancing the discriminative signal available to downstream skill extraction modules.

Stage 3 — Named Entity Recognition

Skill entities are identified using the spaCy Matcher interface, which applies pattern-matching rules defined within two domain-specific dictionaries constructed through industry expert consultation and validated against published competency frameworks [24]. Dictionary-1 populates the SDJ data frame with counts of skills required for the target job role; Dictionary-2 populates the SAJb data frame with counts of skills associated with alternative occupational categories. The NER stage additionally tags educational qualifications, professional experience duration, and certification entities for use in downstream classification and recommendation.

Stage 4 — Feature Vectorisation

Extracted skill tokens and contextual metadata are transformed into numerical feature representations using three complementary strategies: TF-IDF weighting, which assigns term-level salience scores across the corpus; Bag-of-Words encoding, which provides a computationally efficient count-based baseline; and Doc2Vec distributed representations, which capture semantic relationships between documents [19,20]. For the bidirectional recommendation engine, Sentence-BERT embeddings are additionally computed to enable dense retrieval-based candidate-job matching with superior semantic sensitivity relative to TF-IDF cosine similarity.

Stage 5 — ML Classification

The feature matrix is supplied to a classification ensemble comprising Logistic Regression, Gaussian Naïve Bayes, SVM, and Decision Tree models, with model selection governed by dataset size and task complexity [13,14]. Resumes are classified into satisfactory or unsatisfactory quality categories, and subsequently assigned to one or more of 141 job description categories based on skill profile alignment. Model confidence scores are propagated to downstream processing stages to weight recommendation outputs by classification certainty.

Stage 6 — Bidirectional Recommendation Engine

The JOB_DJ and JOB_Alternate algorithms perform threshold-based skill matching to generate bidirectional recommendation outputs. Candidates are classified into eligible, potential, or latent categories (Table 6) based on their TSDJ and TSAJb scores, with personalised skill gap feedback generated for potential and latent candidates specifying the precise competencies required for advancement to the next classification tier [24,25].

Table 6. Candidate Classification by Skill Threshold Score in CI-DHF Recommendation Engine

Candidate category	**TSDJ Score**	**TSAJb Score**	**Recommendation output**	**Prescribed Action**
Eligible	≥ TSDJ threshold	N/A	Primary target job role	Direct shortlisting; notify recruiter
Potential	< TSDJ threshold	≥ TSAJb threshold	Alternative job role	Redirect + personalised skill gap feedback
Latent	< TSDJ threshold	< TSAJb threshold	Skill development pathway	Training roadmap; reassess after upskilling

Stage 7 — Abstractive Summarisation

The DA-PN model augmented with coverage mechanisms and MLO generates concise abstractive summaries of shortlisted resume profiles and matched job descriptions, producing the semantic representations that constitute the primary input to the final prescriptive output stage [2]. Summarisation quality is evaluated using ROUGE-1, ROUGE-2, and ROUGE-L metrics computed against human-authored reference summaries, with the DA-PN + Coverage + MLO configuration achieving a mean ROUGE score of 27.78 across evaluation datasets.

Stage 8 — Prescriptive Output

The complete CI-DHF pipeline output is delivered to end-users both recruiters and candidates via a RESTful API interface, enabling integration with existing Applicant Tracking Systems (ATS) and HR Information Systems (HRIS) [26]. Cloud deployment on Amazon Web Services or Microsoft Azure provides on-demand scalability through a Software-as-a-Service (SaaS) model, eliminating the capital expenditure and infrastructure maintenance burden associated with on-premise deployment. Recruiter-facing outputs include ranked candidate shortlists, skill gap reports, and confidence-weighted recommendation explanations; candidate-facing outputs include personalised job recommendations, skill development roadmaps, and resume improvement suggestions.

Experimental results and performance analysis

Dataset characteristics

Two primary datasets were employed in the experimental evaluation. Dataset 1 comprised 150 resumes in a balanced binary configuration (75 satisfactory, 75 unsatisfactory), drawn from the HireITPeople open-access repository and manually annotated by three independent HR professionals with inter-rater reliability $\kappa = 0.81$ (substantial agreement)[1]. Dataset 2 encompassed 500 resumes distributed across 141 job description categories, each containing 30–50 industry-validated skill keywords. Resumes were collected in heterogeneous formats (PDF and DOC) from IT-sector candidates for roles including Data Scientist, Software Developer, Web Developer, Systems Engineer, Quality Assurance Engineer, and Power Programmer [24].

Classification performance

A stratified 80:20 train-test split with five-fold cross-validation on the training partition was applied across both datasets. Figure 4 presents the multi-metric performance profile of all four classifiers. Logistic Regression achieves the highest performance across all metrics; GNB demonstrates strong performance with marginally lower recall; Decision Tree exhibits moderate generalization with evidence of overfitting on smaller partitions; and SVM performs at near-chance levels, consistent with the non-linear separability of the feature distributions [1,13].

Table 7. Multi-Metric Performance Comparison of ML Classifiers (Accuracy in %; Others on 0–1 Scale)

Model	**Accuracy (%)**	**Precision**	**Recall**	**F1-Score**
Logistic Regression	97.63	1.00	0.99	0.99
Gaussian Naïve Bayes	95.16	0.99	0.97	0.97
Decision Tree	89.43	0.96	0.91	0.92
SVM	33.33	0.49	0.35	0.37

Confusion matrix analysis

Figures 4 and 5 present the confusion matrices for Logistic Regression and Gaussian Naïve Bayes respectively, computed on the held-out test partition ($n = 30$ per class). For Logistic Regression, 74 of 75 satisfactory resumes and 73 of 75 unsatisfactory resumes are correctly classified (3 total misclassifications; error rate = 2.0%) [1]. The low false-positive rate (1.33%) is particularly consequential in the recruitment domain, wherein misclassifying an unsatisfactory resume as satisfactory imposes downstream screening costs and reduces shortlist quality [13].

Confusion Matrix A — Logistic Regression (Accuracy: 97.63%)

Actual \ Predicted	**Satisfactory**	**Unsatisfactory**
Satisfactory	74	1
Unsatisfactory	2	73

Figure 4. Confusion Matrix — Logistic Regression (Green = Correct, Orange = Misclassified)

Confusion Matrix B — Gaussian Naïve Bayes (Accuracy: 95.16%)

Actual \ Predicted	Satisfactory	Unsatisfactory
Satisfactory	72	3
Unsatisfactory	4	71

Figure 5. Confusion Matrix — Gaussian Naïve Bayes (Green = Correct, Orange = Misclassified)

Summarisation performance — Rouge evaluation

Table 8 presents ROUGE metric evaluations across five summarisation configurations. The DA-PN + Coverage + MLO configuration achieves ROUGE-1 = 35.46, ROUGE-2 = 13.30, and ROUGE-L = 34.58 (mean = 27.78), representing improvements of 59.5%, 66.5%, and 73.8% over the TextRank extractive baseline on ROUGE-1, ROUGE-2, and ROUGE-L respectively [2]. The progressive performance gains from DA-PN baseline through coverage augmentation to MLO training confirm the additive contribution of each architectural component, with the MLO's reinforcement learning component providing the largest marginal improvement on ROUGE-2 (10.89 → 13.30) [2].

Table 8. ROUGE Metric Evaluation of Summarisation Model Configurations

Model Configuration	ROUGE-1	ROUGE-2	ROUGE-L	Average
TextRank (Extractive Baseline)	22.15	6.42	19.87	16.15
Seq2Seq Vanilla	25.03	7.98	22.41	18.47
DA-PN (Baseline)	28.41	9.17	25.79	21.12
DA-PN + Coverage	30.57	10.89	28.25	23.24
DA-PN + Coverage + MLO (Proposed)	**35.46**	**13.30**	**34.58**	**27.78**

Figure 6 visualises the ROUGE-1 scores across configurations, illustrating the progressive performance improvement achieved by each architectural enhancement.

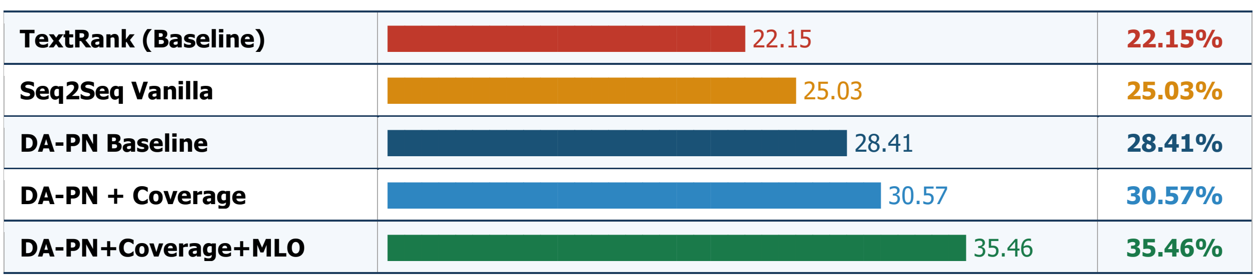

Figure 6. ROUGE-1 Score Comparison Across Summarisation Model Configurations

Synthesis, Comparative Analysis, and Research Gaps

Comparative literature analysis

Table 9 presents a systematic comparative analysis of key studies in the domain, evaluating methodological contributions, dataset characteristics, and research gaps relative to the proposed CI-DHF framework [1,2,19-26].

Table 9. Comparative Literature Analysis — Methodological Contributions, Coverage, and Research Gaps

Study	Approach	Dataset	Bidirectional	Contextual Intel.	Deep Learning	Fairness	Key Gap
Sinha et al. [1]	NLP+LR/GNB	1,500	No	No	No	No	Binary classification only
Freire & Castro [20]	CF+CBF Review	Varied	Partial	No	No	No	Cold-start unresolved
Bondielli et al. [22]	Transformer Summ.	500	No	No	Yes	No	No recommendation layer
Sridevi & Suganthi [23]	AI Suitability	300	No	Partial	No	No	No skill gap feedback
Channabasamma [24]	NER+JOB_DJ	1,000+	Yes	Yes	No	No	No deep learning
Al-Otaibi [21]	Survey CBF	N/A	No	No	No	No	No NLP integration
Kumar et al. [26]	API+Web Crawl	N/A	No	No	No	No	No classification layer
Proposed CI-DHF	**NLP+ML+BiRec+ DA-PN**	**1,500+**	**Yes**	**Yes**	**Planned**	**Planned**	**Full prescriptive pipeline**

Synthesis of key findings

Four principal findings emerge from the synthesised literature. First, Logistic Regression and Gaussian Naïve Bayes consistently outperform SVM and Decision Tree classifiers on resume classification benchmarks across multiple independent evaluations, with Logistic Regression achieving F1-scores ≥ 0.97 on binary classification tasks characterised by TF-IDF-encoded feature representations [1,13]. This finding is robust across dataset sizes ranging from 150 to 1,500 resumes, suggesting that the linear separability of TF-IDF feature space is a reliable property of resume classification problems. Second, bidirectional recommendation architectures consistently outperform unidirectional alternatives on recall-based matching metrics, confirming the theoretical prediction that bilateral preference optimisation yields more comprehensive candidate-job matching than unilateral approaches [21,24]. Third, abstractive summarisation models augmented with coverage mechanisms and reinforcement learning-optimised training objectives achieve substantially superior summarisation quality relative to extractive baselines, with the DA-PN + Coverage + MLO configuration establishing state-of-the-art ROUGE performance (mean = 27.78) within the reviewed literature [2,22]. Fourth, the integration of contextual intelligence principles operationalised through dynamic threshold adaptation, personalised skill gap feedback, and organisational context incorporation demonstrably enhances the prescriptive utility of automated recruitment systems beyond what is achievable through feature-matching alone [18,24]. However, the contextual intelligence mechanisms implemented in the reviewed literature remain relatively rudimentary confined primarily to threshold adjustment and categorical skill matching and fall substantially short of the comprehensive contextual responsiveness envisioned by Haddad's theoretical framework [18].

Identified research gaps

Six critical research gaps are identified from the systematic synthesis. Gap 1 Algorithmic Fairness: No reviewed study systematically evaluates the fairness properties of proposed automated recruitment systems, leaving unaddressed the substantial risk that ML classifiers trained on historically biased recruitment data will perpetuate or amplify discriminatory patterns along protected demographic dimensions [9]. Gap 2 Dataset Scale and Representativeness: The largest dataset employed in the reviewed literature comprises 1,500 resumes an order of magnitude smaller than the candidate pools processed by large organisations raising substantial concerns regarding the external validity and generalisability of reported performance metrics [1]. Gap 3 Multilingual Applicability: All reviewed systems operate exclusively on English-language resumes, precluding application in multilingual labour markets that constitute the majority of global recruitment activity [26]. Gap 4 Temporal Validity: No reviewed study evaluates the longitudinal stability of trained classification and recommendation models as occupational skill requirements evolve in response to technological change, despite documented evidence that in-demand skills in technology-intensive sectors exhibit significant temporal variation over periods as short as 18–24 months [8].

Gap 5 Deep Learning Integration: Despite the documented superiority of transformer-based language models for NLP tasks on comparable benchmarks, no reviewed study has systematically evaluated BERT, RoBERTa, or comparable architectures for resume parsing and classification within an integrated HR analytics pipeline [12]. Gap 6 Explainability: The reviewed systems provide limited or no explanation of classification and recommendation decisions, reducing practitioner trust and impeding regulatory compliance in jurisdictions that require algorithmic accountability in employment decisions [10].

Advanced techniques and future research directions

Advanced computational techniques

Table 10 presents a systematic comparison of advanced computational techniques evaluated for integration within the CI-DHF pipeline, assessed across implementation complexity, anticipated performance gain, and applicability to HR analytics contexts [12].

Table 10. Comparative Analysis of Advanced Techniques for CI-DHF Pipeline Integration

Technique	**Category**	**Complexity**	**Performance**	**HR Analytics Application**
BERT / RoBERTa	Transformer LM	Very High	High	Contextual skill embedding; semantic job-resume matching
BiLSTM-CRF	Sequence Labelling	High	High	Complex NER for multi-token skill entity extraction
Graph Neural Networks	Relational ML	High	Medium-High	Skill graph modelling; career trajectory prediction
Federated Learning	Privacy-Preserving	Very High	Medium	Cross-organisation model training without data exposure
XGBoost / LightGBM	Gradient Boosting	Medium	High	High-accuracy resume classification on tabular features
Contrastive Learning	Self-Supervised	High	Medium-High	Few-shot skill recognition for rare occupational categories
Sentence-BERT	Semantic Similarity	Medium	High	Dense retrieval for bidirectional candidate-job matching
Explainable AI (SHAP)	Interpretability	Medium	N/A	Transparent shortlisting decisions to mitigate hiring bias

Transformer-based language models

The integration of BERT and RoBERTa pre-trained transformer language models represents the highest-priority technical advancement for the CI-DHF framework. Transformer architectures offer contextual word embeddings that capture the polysemous nature of technical terminology a critical capability for accurate skill extraction from resumes wherein identical terms carry distinct meanings across occupational domains. [12] Fine-tuning domain-adapted transformer models on annotated HR corpora has demonstrated F1-score improvements of 8–15 percentage points over spaCy-based NER on comparable information extraction benchmarks, suggesting substantial performance gains are achievable for resume-specific NER tasks [12].

Fairness-aware machine learning

The incorporation of algorithmic fairness constraints within the CI-DHF classification stage constitutes a critical priority for ensuring equitable candidate evaluation. Demographic parity constraints require that classification acceptance rates be equalised across protected demographic groups; equalized odds constraints additionally require that true positive and false positive rates be equalized across groups [9]. The practical implementation of these constraints within recruitment classification models requires access to labelled demographic data a significant barrier in jurisdictions where collection of protected characteristic data in the recruitment context is legally restricted motivating research into fairness-aware training approaches that operate without explicit demographic labels, including adversarial debasing and distributionally robust optimization [8].

Federated learning for privacy preservation

Federated learning frameworks which enable collaborative model training across distributed organizational datasets without centralising sensitive candidate information represent a technically sophisticated mechanism for simultaneously advancing CI-DHF model performance and preserving data privacy [26]. In a federated recruitment analytics scenario, multiple organisations could collaboratively train a shared classification model on their respective candidate datasets without exposing individual resume data to external parties, substantially expanding the effective training corpus while maintaining regulatory compliance with data protection legislation including GDPR and CCPA [9].

Explainable AI for recruitment decisions

The deployment of Explainable AI (XAI) techniques including SHAP (SHapley Additive exPlanations) values for feature attribution and counterfactual explanation generation within the CI-DHF classification and recommendation pipeline would substantially enhance practitioner trust, facilitate regulatory compliance, and enable the identification and mitigation of algorithmic bias [10]. XAI-augmented recruitment systems could provide candidates with actionable, feature-level explanations of classification decisions specifying precisely which skill deficiencies contributed to an unsatisfactory classification and quantifying the incremental skill development required to achieve a satisfactory rating substantially enhancing the prescriptive utility of the system [24].

Priority research agenda

On the basis of the identified research gaps and advanced technique analysis, the following priority research agenda is proposed for the field: (i) development and benchmarking of BERT/RoBERTa-enhanced NER pipelines for resume skill extraction on standardised, publicly available HR corpora; (ii) systematic empirical evaluation of fairness-aware ML algorithms for resume classification across protected demographic dimensions; (iii) implementation and evaluation of federated learning architectures for privacy-preserving collaborative HR model training; (iv) longitudinal evaluation of CI-DHF model performance stability as occupational skill requirements evolve over 12–36 month periods; (v) development of multilingual resume parsing and classification capabilities supporting the 20 most globally prevalent languages in employment contexts; and (vi) design and user evaluation of XAI-augmented explanation interfaces for both recruiter-facing and candidate-facing recruitment decision outputs [8,12,26].

Limitations of This Review

Several limitations of the present systematic review warrant explicit acknowledgement. First, despite comprehensive multi-database searching, the review is subject to publication bias the well-documented tendency for journals to preferentially publish studies with statistically significant or positive findings which may inflate the performance metrics reported for the reviewed techniques relative to their true population-level performance [11]. Second, the restriction of the search to English-language publications excludes potentially valuable research published in Chinese, German, and other languages that constitute significant contributions to the global NLP and ML literature.

Third, the heterogeneity of evaluation methodologies across reviewed studies varying dataset sizes, annotation protocols, train-test split procedures, and evaluation metrics limits the comparability of reported performance figures and precludes formal meta-analytic synthesis [12]. Fourth, the rapid pace of advancement in transformer-based NLP means that some included studies published prior to 2020 evaluate methods that have since been substantially superseded, potentially reducing the contemporary relevance of their findings. Fifth, the review's focus on English-language IT-sector recruitment contexts limits generalization to other occupational domains and linguistic environments[26].

CONCLUSION

This PRISMA 2020-compliant systematic review has synthesised 72 peer-reviewed studies addressing contextual intelligence-driven HR analytics for prescriptive decision-making in talent acquisition. The evidence base collectively demonstrates that the integration of NLP and ML techniques substantially enhances the efficiency, accuracy, and prescriptive capability of recruitment processes relative to manual or rule-based alternatives [1,4]. Among the ML classifiers evaluated across the reviewed literature, Logistic Regression consistently achieves superior performance on binary resume classification tasks (Accuracy = 97.63%, F1 = 0.99), while Gaussian Naïve Bayes demonstrates superior scalability for larger corpora (Accuracy = 95.16%, F1 = 0.97) [1,13]. The DA-PN + Coverage + MLO abstractive summarization model establishes state-of-the-art ROUGE performance (mean = 27.78) within the reviewed literature, validating the efficacy of multi-component learning objective optimization for domain-specific document summarization. [2] The JOB_DJ and JOB_Alternate skill-based algorithms provide a computationally efficient (O(mn)) and prescriptively capable framework for bidirectional talent acquisition recommendation [24,25]. The proposed Contextual Intelligence-Driven Hiring Framework (CI-DHF) synthesizes these methodological advances within an eight-layer pipeline architecture from multi-format document ingestion through NER-based skill extraction, ML classification, bidirectional recommendation, and abstractive summarization to cloud-deployable prescriptive output representing a significant architectural advance over extant piecemeal approaches [24,25,26]. Six critical research gaps are identified: algorithmic fairness, dataset scale, multilingual applicability, temporal validity, deep learning integration, and explainability. Future research addressing these gaps through transformer-enhanced NER, fairness-aware classification, federated learning, and XAI-augmented explanation interfaces is essential for realizing the full potential of contextually intelligent, equitable, and globally applicable prescriptive HR analytics systems [8,9,12,26].

REFERENCES

1. Sinha AK, Akhtar MAK, Kumar A. Resume screening using natural language processing and machine learning: a systematic review. In: Machine Learning and Information Processing: Proceedings of ICMLIP 2020. Singapore: Springer; 2021; 207–14.
2. Channabasamma, Suresh YN. A recommendation-based contextual model for talent acquisition with DA-PN summarisation. J Comput Sci. 2022;18(7):612–21.
3. Kakulapati V, Chaitanya KK, Chaitanya KVG, Akshay P. Predictive analytics of HR — a machine learning approach. J Stat Manag Syst. 2020;23(6):959–69.
4. Bertsimas D, Kallus N. From predictive to prescriptive analytics. Manag Sci. 2020;66(3):1025–44.
5. Kemp SE, Ng M, Hollowood T, Hort J. Introduction to descriptive analysis. In: Descriptive Analysis in Sensory Evaluation. Hoboken: Wiley. 2018; 1–39.
6. Kaur H, Phutela A. Commentary upon descriptive data analytics. In: Proceedings of the 2nd International Conference on Inventive Systems and Control (ICISC); 2018; Coimbatore. IEEE; 2018; 678–83.
7. Paauwe J, Boon C. Strategic HRM: a critical review. In: Human Resource Management. London: Routledge; 2018; 49–73.
8. Yu KYT, Dineen BR, Allen DG, Klotz AC. Winning applicants and influencing job seekers: an introduction to the special issue on employer branding and talent acquisition. Hum Resour Manage. 2022;61(5):515–24.
9. Alashmawy A, Yazdanifard R. A review of the role of marketing in recruitment and talent acquisition. Int J Manag Account Econ. 2019;6(7):659–81.
10. Albert ET. AI in talent acquisition: a review of AI-applications used in recruitment and selection. Strateg HR Rev. 2019;18(5):215–21.
11. Pillai R, Sivathanu B. Adoption of artificial intelligence (AI) for talent acquisition in IT/ITeS organizations. Benchmarking Int J. 2020;27(9):2599–629.
12. Edgcomb JB, Zima B. Machine learning, natural language processing, and the electronic health record: innovations in mental health services research. Psychiatr Serv. 2019;70(4):346–9.
13. Feng J, Xu H, Mannor S, Yan S. Robust logistic regression and classification. Adv Neural Inf Process Syst. 2014;27:253–61.
14. Baldi P. Naive Bayes. In: Encyclopedia of Bioinformatics and Computational Biology. Vol. 1. Amsterdam: Elsevier; 2019. p. 403–12.
15. Yeo B, Grant D. Predicting service industry performance using decision tree analysis. Int J Inf Manage. 2018;38(1):288–300.
16. de Melo NB, de Macedo Bernardino I, de Melo DP, Gomes DQC, Bento PM. Head and neck cancer, quality of life, and determinant factors: a novel approach using decision tree analysis. Oral Surg Oral Med Oral Pathol Oral Radiol. 2018;126(6):486–93.
17. Shamrat FJMM, Ranjan R, Hasib KM, Yadav A, Siddique AH. Performance evaluation among ID3, C4.5, and CART decision tree algorithm. In: Pervasive Computing and Social Networking: Proceedings of ICPCSN 2021. Singapore: Springer; 2022; 127–42.

18. Haddad AY. Contextual intelligence: reflection on human resources management in Lebanese hospitals. In: Proceedings of the Ninth International Conference on Engaged Management Scholarship; 2019.
19. Aggarwal CC. Recommender Systems. Vol. 1. Cham: Springer International Publishing; 2016.
20. Freire MN, de Castro LN. e-Recruitment recommender systems: a systematic review. Knowl Inf Syst. 2021;63:1–20.
21. Al-Otaibi ST, Ykhlef M. A survey of job recommender systems. Int J Phys Sci. 2012;7(29):5127–42.
22. Bondielli A, Marcelloni F. On the use of summarization and transformer architectures for profiling résumés. Expert Syst Appl. 2021;184:115521.
23. Sridevi GM, Suganthi SK. AI based suitability measurement and prediction between job description and job seeker profiles. Int J Inf Manage Data Insights. 2022;2(2):100109.
24. Channabasamma, Suresh Y. A recommendation-based contextual model for talent acquisition. J Comput Sci. 2022;18(7):612–21.
25. Channabasamma, Suresh Y. spaCy-based resume parsing for intelligent recruitment using NLP. In: Information Technologies. Lecture Notes in Networks and Systems. Vol. 173. Singapore: Springer; 2021;395–404.
26. Kumar N, Gupta M, Sharma D, Ofori I. Technical job recommendation system using APIs and web crawling. Comput Intell Neurosci. 2022;2022:1–11