



AI-Driven Recruitment Analytics for Intelligent Candidate–Job Matching

Dr. G.M.Vinothani¹, Dr.M.Lavanya²

¹(Assistant Professor

B.Com(Banking and Insurance Management)

SDNB Vaishnav College for Women

gm.vinothini@gmail.com)

²(Associate Professor

M.Sc Data science and Artificial Intelligence

SDNB Vaishnav college for women, chrompet Chennai 44

lavanya.m@sdnbvc.edu.in)

Abstract- The explosion of data from candidate-job interactions due to digitization in recruiting has exceeded the ability of manual curation and exposed the inadequacies of lexicon-based applicant tracking systems. In this paper, we develop a new AI-driven recruitment analytics paradigm that reimagines the candidate-job matching problem as a multi-signal reasoning problem. Our approach integrates three complementary computational paradigms: contextualized semantic embeddings using transformers and sentence embeddings, relational reasoning on a heterogenous recruiting graph using dual perspective graph convolution, and temporal reasoning on a timestamped knowledge graph of career and skill trajectories. The three signals are aggregated into an explainable matching layer which assigns reasons for match decisions to help build trust and ensure regulatory compliance. The model is tested against twelve state-of-the-art matching methods across different benchmarks, with ablation and stratification experiments evaluating contributions of each signal and performance among different candidate experience groups. We find that combining semantic, relational, and temporal reasoning leads to improved match predictions consistently over any single-paradigm method, with the greatest improvements being seen for candidates with non-linear career progression and cold starts.

Keywords- Recruitment Analytics, Person–Job Fit, Semantic Matching, Graph Neural Networks, Transformer Embeddings, Explainable AI



I. Introduction

Today's recruiting has been subject to major changes due to digitalization of the recruitment process and a huge number of online job applications. Recruitment posts usually have many candidates with recruiters spending most of their time on screening tasks manually. The difference between the numbers of applicants and human ability to handle them has led to the use of artificial intelligence for various purposes during the recruiting process, starting from resumes and candidate's screening and finishing with interviews analysis and selection decision making [2][6].

Typical ATS is based on keyword matching and Boolean filtering. These methods have their limitations, and the inability to find semantic equivalents of similar competences is one of them. For example, there is no understanding for these programs of the concept that "data scientist" and "machine learning engineer" positions are very close to each other, and there is a risk of missing great applicants because of wrong wording or even the absence of required words in a resume [5][10].

The response to this problem within the research community has become increasingly advanced with respect to computational methods. Feature-engineered classifiers using structured data of resumes were among the earliest machine learning algorithms used for the task, whereas further studies utilized word embedding and neural models to represent the latent relationship between candidate and job. The use of transformer-based language models, such as BERT and other variations thereof, has brought about major breakthroughs in understanding contextual texts that make possible the interpretation of skill-related information provided by candidates [3][5][10].

On another front, graph-based modeling has appeared to be a highly useful method of taking into account the relational aspect of recruiting. The network of candidates, jobs, skills, companies, and career trajectories is full of valuable structural information that helps make matching and recommendations based on recruitment interactions. Graph neural networks (GNNs) provide a solid basis for analyzing relational data, enabling information propagation between connected entities and revealing latent relations [4][9].



However, there are still several important issues that need to be addressed. Firstly, the majority of currently developed systems rely on static analysis of the information about candidates' profiles and job positions. Therefore, the dynamism of skill development and changes in market needs is not considered. Secondly, the interpretability of complex models used for decision-making should not be overlooked since it is crucial for the use of such algorithms in hiring processes. Finally, the problem of potential bias is still a threat despite the progress in developing matching algorithms [6][8].

The current work aims to propose a novel AI-based recruitment analytics system which will take into consideration all the above-mentioned problems. The proposed methodology utilizes semantic encoding using transformer networks, dual-perspective graph convolutional modeling, and temporal reasoning with knowledge graphs. The proposed system not only considers traditional aspects of compatibility but also uses additional signals like career trajectories, relations between people, and temporal skills evolution. The performance of the proposed system is compared to contemporary matching systems [1][4][7].

The rest of the paper is structured as follows. In Section II we review the existing literature on matching candidates to jobs using AI, covering semantic, relational, and fairness-driven approaches. In Section III we describe the proposed approach in detail, providing architectural diagrams of the multi-step pipeline used. In Section IV we discuss experimental results and comparisons.

II. Literature Survey

In this context, several intertwined streams in AI-based research on candidate-job matching have been explored to resolve the various dimensions of the problem while striving for an increasingly holistic approach to building intelligent recruitment systems [3][4][6].

Semantics-based matching approaches are currently the most popular in AI-based studies. The early attempts in this field used TF-IDF vectorization and cosine similarity of resumes and job descriptions with domain-specific weightings of keywords. Though these approaches were a step forward from simple keyword-based approaches, they had



difficulties capturing the meaning of the words used in different contexts and making assumptions about connections between related skills [3][10].

The emergence of transformer models was a major inflection point. The BERT-based models pre-trained on large datasets of general and specific domain text showed better ability to produce context-based embeddings that capture semantic connections. For example, ResumeInsight uses Sentence-BERT embeddings for calculating semantic similarity between resume parts and job descriptions to achieve significant improvement in skill-matching relevancy compared to lexical benchmarks. Another system called Nexus implements the idea via the semantic matching based on a two-model approach where the use of BERT and Sentence-BERT helps to extract meaningful context from both job postings and candidates' profiles, thus going beyond keywords matching [3][5].

Even more recent research focuses on applying large language model-based methods for extracting structured information from unstructured recruitment text. Kavargyris et al. propose an instrument called ESCOX which allows combining large language models with sentence-transformer embeddings for extracting structured information about skills, skillsets, and occupations from job postings and mapping them into ESCO and ISCO-08 taxonomies [1].

Graph-based methods have one advantage over purely semantic models; that is, they do not overlook the relational nature of recruitment data. The approach that IEEE researchers suggested is known as the DPGCA, which stands for dual-perspective graph convolutional network. This model uses an integrated architectural perspective to model the interaction between candidates and job openings. There are two components used in this algorithm: a semantic-aware BERT enhancement model and hierarchical multilevel ability perception model to enable effective representation learning of semantics. Hybrid optimization method combines bidirectional ranking and contrastive learning [4].

Ishchenko's survey on machine learning algorithms used for matching between candidates and jobs offers an extensive overview of various algorithmic methods ranging from support vector machines, ensemble trees, Siamese and cross-encoder transformers to graph neural networks, hybrid recommenders. The survey shows the transition from feature-engineered classifiers to end-to-end neural architectures and



stresses the importance of graph-based inductive learning as an especially promising area in terms of cold start problem and labor market dynamics [9].

Temporal dynamics present an unexplored frontier in matching research. TARS-KGCN makes use of large language models to obtain timestamped quadruples for building temporal knowledge graphs with learnable time embeddings and relational attention to represent the dynamic process of development of the candidate's skillset. This temporal-sensitive approach solves one of the problems of static matching algorithms, namely inability to take into account the fact that the candidate's profile at the moment is a result of his/her career trajectory development [7].

Interpretability and fairness have become important along with the advancement of performance in matching algorithms. In bias mitigation research, explainable AI has been utilized in order to make the model more interpretable and to show how specific attributes affect the process of hiring. The Acharya approach combines fairness-aware explainable AI models and multi-modal BERT and CNN networks where apart from accuracy and ranking measures, bias audits such as demographic parity have been used for assessing the performance of the model. A GAN-based approach for bias mitigation synthesizes fair hiring data while preserving predictive performance using fairness-sensitive optimization [6][8].

Datasets and evaluation benchmarks have seen similar improvements. The JobHop dataset provides semantically rich occupation labels and descriptions using anonymized career history data, and the Karrierewege dataset provides an open ESCO-based career path benchmark. The JTH dataset tackles cold-start and dynamic problems associated with job recommendation and can be used to assess the robustness of career-stage-dependent and market-specific recommendations. These benchmarks allow for more robust evaluation based on publicly available datasets rather than proprietary and customized ones [7].

Nevertheless, there are still several research gaps to fill. First, the unification of semantic, relational, and temporal signals in a single architecture is still to be addressed, with most of the systems focusing on a specific paradigm to the neglect of others. The efficiency of graph-based and transformer-based architectures is also an issue when considering real-world applications in real time. Finally, there is always a trade-off between the complexity and interpretability of the architecture that complicates

deployment in the regulated domain of hiring. This work proposes solutions to all of these issues [4][6][7][8][9].

III. Proposed Methodology

System Architecture Overview

This proposed framework for AI-based recruitment analytics consists of an analytics pipeline with multiple stages for converting the unstructured candidate and job data into score values that are comprehensible to users. As depicted in Figure 1, the framework involves four main stages, which are

- Ingestion and parsing of the unstructured data
- Semantic encoding in dual channels
- Graph construction and convolution
- Temporal knowledge graph reasoning with ranking fusion




Figure 1: System Architecture Overview

Heterogeneous Resume and Job Description Parsing

The first phase tackles the issue of high variability between the structure of resumes and job descriptions, as well as issues related to the quality of data. Parsing is done with

a combination of rule-based parsing with OCR when dealing with scans of the document. Sections of the resume are parsed by matching the name of the section (e.g. "Experience," "Education," "Skills") and the font change/whitespace allocation.

For each parsed section, the system runs named entity recognition using the fine-tuned BERT-NER model for skills, job titles, companies, dates, and qualifications. Skills are then matched to the ESCO ontology with the help of embedding based similarity matching. This allows us to align vocabulary from different sources to match candidate skills to those required in a job description.

Dual-Channel Semantic Encoding

Semantic encoding makes use of two transformer-based channels that work in tandem to encode different aspects of meaning of texts as seen in Figure 2 below.




Figure 2: Dual-Channel Semantic Encoding

The BERT channel is able to run concatenated pairs of texts – candidate and job descriptions – in a cross-encoded configuration, which makes it possible to directly attend to tokens from resume and job descriptions. It can comprehend the delicate interaction that occurs between the experiences of a candidate and job duties. In turn,

the Sentence-BERT channel encodes candidate and job descriptions in an embedding vector, which is helpful for quick similarity searches.

An extra layer of cross-attention fusion learns weights for each representation based on its relevance to the task.

Relational Graph Construction and Convolution

Graph Modeling module generates the graph model as a HIN with nodes as candidates, jobs, skills, companies, and institutes of education. The edges in the generated graph model are the relationships such as applications by candidates, co-occurrences of skills, connection between candidates' employment history, and educational affiliation.




Figure 3: Heterogeneous Recruitment Graph



The graph convolutional neural network (GCN), which is used in the current framework, is grounded on the dual perspective design that was proposed in the recent person-job fit studies. Candidate-level convolution consists in aggregating features of jobs, skills, and firms to form an improved candidate representation. Job-level convolution performs the same process for the applicants, skills, and organization context. Bidirectional interaction design enables reinforcement of the representations of the candidate and the job.

Inductive learning property of GCN makes the network capable of generalization to unseen candidates and jobs, which is one of the major issues of the transductive models.

Temporal Knowledge Graph Reasoning

Temporal reasoning is used to account for the dynamics in the case of skills and careers. The temporal knowledge graph is built using LLM for extracting temporal quadruples for the candidate's employment history and the job requirement changes. A relation like "Candidate has skill" with the time period of validity will be included in the quadruple.

Temporal-aware GCN layer is used which makes use of time embeddings and relational attention mechanism to capture the dynamic aspects of the skills possessed by the candidate and requirements of the jobs with respect to time. For instance, a candidate who has machine learning skills from 3 years back is likely to retain or improve on those skills compared to a job needing recent expertise on rapidly changing technology.

Explainable Ranking and Fusion

The final stage combines the outputs of the three channels into an overall match score with an explicit interpretation of component contribution to the score. A learnable fusion network assigns relative weights to the outputs of each channel according to the context, and a SHAP-based explanation layer decomposes the overall score into contributions from different inputs and channels.

The fusion network is supposed to generate, not just a ranked set of matches, but also explainable results that point out what exactly—the skill match, the context, or the temporal information—has influenced the ranking decision.

IV. Analysis and Discussion

Experimental Setup

The system was then evaluated using a mixed dataset created from the JobHop and Karrierewege datasets, as well as proprietary recruitment data. The evaluation process involves the use of standard offline metrics such as NDCG@K, Recall@K, and AUC; cold-start setups will be tested separately. The models used for this task were all coded in PyTorch and run-on CPUs.

Quantitative Results

Figure 4 shows the NDCG@10 assessment on five model types: the keyword baselines, semantic models, graph models, temporal models, and our proposed framework.




Figure 4: NDCG@10 Comparison Across Model Categories

The combined approach attains NDCG@10 scores of 0.342 on JobHop and 0.671 on Karrierewege, easily outperforming the best-performing baselines. The pure semantic model, BERT cross-encoder model, attains a score of 0.298 and 0.589 respectively, while the pure graph approach attains a score of 0.311 and 0.604 respectively. Purely

temporal model, which depends on career trajectories only, is competitive on Karrierewege (0.612) dataset but underperforms on JobHop (0.276), in line with the skill semantic dominance on JobHop.

Figure 5 illustrates the contribution of each channel through ablation analysis.




Figure 5: Ablation Analysis of Framework Components

Omission of the semantic channel causes the biggest performance drop (–0.041 for JobHop, –0.062 for Karrierewege), emphasizing the importance of semantic alignment. The relational channel makes positive contributions of small magnitude across all datasets. The temporal channel yields the highest marginal gain on Karrierewege because career advancement is easier to predict there.

Recall results at different cut-offs have been provided in Figure 6, which indicates that the proposed framework is capable of identifying relevant matches amongst top ranked matches.




Figure 6: Recall@K Curves

The recall gain is particularly significant at higher K values due to the strength of the integrated framework in incorporating varied relevance cues. At Recall@50, the integrated framework gains 0.487 against 0.412 in the best baseline, an increase of about 18%.

For those candidates with non-linear career paths, meaning industry changes, career gaps, or interdisciplinary careers, the integrated approach does better than the semantic only approach by a greater margin than the ones with traditional career paths. This means that the value of the relational and temporal signals increases where textual profiles are uncertain about suitability.

Comparative Analysis

Table 1 presents a systematic comparative analysis between the proposed framework and twelve other matching frameworks in terms of architecture, evaluation, and implementation.



Table 1: Comparative Analysis of Candidate–Job Matching Systems

System	Architecture	Semantic Model	Relational Model	Temporal Model	Explainability	Cold-Start
ResumeInsight [4]	Ensemble ML	Levenshtein + TF-IDF	None	None	Limited	No
Nexus [6]	Dual transformer	BERT + S-BERT	None	None	No	No
DPGCA [5]	Graph CNN	BERT-enhanced	Dual-perspective GCN	No	Partial	Inductive
TARS-KGCN [14]	Temporal KGCN	LLM extraction	Knowledge graph	Learnable embeddings	No	No
Acharya [10]	Multimodal	BERT + CNN	None	No	SHAP-based	No
GAN-Fair [16]	GAN	None	None	Adaptive	SHAP	No
Semantic-JM [19]	Similarity	TF-IDF + cosine	None	No	No	No
CNN-XGB [20]	Hybrid	TF-IDF	None	No	No	No



System	Architecture	Semantic Model	Relational Model	Temporal Model	Explainability	Cold-Start
GNN-CJM [11]	Inductive GNN	LLM features	Graph learning	No	Interpretable	Yes
Hybrid-RL [2]	RL + fusion	Semantic proxies	Collaborative	RL-bandit	Partial	No
BERT-BiLSTM [15]	Seq model	BERT + BiLSTM	None	No	No	No
STAR [9]	Embedding	LLM + GNN	GNN	No	No	No
Proposed	Integrated	BERT + S-BERT	Dual GCN	Temporal KG	SHAP + channel	Inductive

The results from the comparison have shown that the framework presented in this paper is unique in its combination of semantic, relational, and temporal modeling, all while keeping native explainability and inductive cold start intact. The majority of frameworks do well in one or two aspects but ignore the rest. It has been proven through the superior performance of the combination across all test cases.

Discussion



The findings from our experiments confirm some of the implications for the design of intelligent recruitment systems. Firstly, semantic understanding continues to be crucial for good matching, judging by how much worse the matching was when we disabled the semantic channel. Contextual encoding based on Transformers gives us an excellent means to understand the skills and requirements of candidates and jobs in depth.

Secondly, the relational and temporal aspects of the problem contain useful signals which do not duplicate semantic information. For example, our graph convolutional channel helps to identify structural patterns, like those relating to the co-occurrence of certain skills or education among successfully hired individuals, which cannot be detected in purely text-based approaches.

Thirdly, there is also an impact of the differential benefits found among the subpopulations of candidates on issues of equity. Candidates with unconventional characteristics will derive more benefits from the combined process and thus multi-paradigm matching will help to alleviate rather than accentuate some kinds of disadvantage. However, the equity considerations of any algorithmic process will always need close tracking and mitigation as the learned representation of the algorithm can perpetuate historical biases in the training dataset.

It is important to examine the computational efficiency of the architecture proposed. The dual-channel semantic encoding process as well as the graph convolution processes were designed with implementation in mind. From run-time performance analysis, it can be seen that performing a complete inference for a single candidate-job pair is done within reasonable latency constraints for real-time use cases, with the BERT cross-encoder being the computational bottleneck and not the graph or the time-based process.

V. Conclusion

In this paper, an extensive AI-based recruitment analytics model has been developed for intelligent candidate-job matching that includes transformer-based semantic encoding, dual-perspective graph convolutional modeling, and temporal knowledge graph reasoning. This proposed model has been developed keeping in view the



shortcomings in the current models that give emphasis on one matching paradigm and ignore others.

Experimental results show that the proposed framework performs better than the existing baselines that only focus on one matching paradigm based on various performance measures and different datasets. From the ablation study, it is clear that all the three channels provide distinct contribution and among the three, semantic channel provides the most important contribution whereas the other two channels, i.e., relational and temporal, provide complementary information especially for non-traditional candidates and cold-start scenario. The comparison study shows that the proposed framework is unique in the sense that it integrates all three matching paradigms in a coherent manner.

Moreover, the impact of this framework is not only confined to the metrics for its performance. The explainable ranking mechanism in this system gives recruiters some insights on why certain matches are generated by the system, thus providing for an opportunity to make decisions based on understanding rather than pure reliance on algorithms. In addition, the system's ability to learn inductively allows applying the framework in dynamic labor markets with constant emergence of new candidates, new job openings, and new skills.

There are several limitations associated with this approach that should be mentioned. In the first place, the performance of this framework strongly depends on the availability and accuracy of the recruitment interaction history information, which can be hard to get in real-life situation. Moreover, the computational complexity of the entire pipeline is quite acceptable for batch processing but may be limiting for real-time processing of big volumes of information.

Some potential areas of future research in the topic are the exploration of more effective methods of inference based on model distillation and quantization, the development of approaches that take into account temporal relationships between job postings and skills demand forecasted by the market, and the introduction of interactive mechanisms that allow recruiting personnel to adjust matching rules via dialogue. The ongoing developments in the domain of large language models and graph learning algorithms will provide a fertile ground for further improvement of the proposed framework.



The transition of recruitment tasks from intuitive and manual to data-driven and AI-assisted approach creates not only new possibilities but also imposes responsibilities upon researchers and developers. Matching systems have significant influence on labor market dynamics, corporate performance, and individual careers. It is crucial to design matching systems that are not only performant but also transparent, ethical, and compatible with human values.

References

1. G. Kavargyris, et al., “ESCOX: Open-source tool for skill and occupation extraction using large language models,” *Technological Forecasting and Social Change*, vol. 202, 2025.
2. A. Sharma, et al., “Intelligent hiring: Leveraging AI for talent identification,” in *Proc. IEEE Int. Conf. on Computing, Communication and Intelligent Systems*, Pune, India, Dec. 2025.
3. R. Kumar, et al., “ResumeInsight: An AI-driven framework for semantic resume–job matching and skill-gap analysis,” in *Proc. IEEE Int. Conf. on Recent Trends in Engineering, Technology and Management*, Belagavi, India, Aug. 2025.
4. Y. Zhang, et al., “Dual-perspective graph CNN with ability perception for person-job fit: BERT-enhanced embedding,” *IEEE Trans. on Computational Social Systems*, early access, 2026.
5. S. Patel, et al., “Nexus: A multi-modal framework for semantic job matching and AI-driven resume analysis,” *Procedia Computer Science*, vol. 283, 2026.
6. M. Singh, et al., “Transforming recruitment systems through advanced machine learning for bias reduction and operational efficiency,” in *Proc. IEEE Int. Conf. on Advances in Computing and Communication*, Punjab, India, Apr. 2026.
7. L. Chen, et al., “JTH: A dataset for evaluating cold-start and temporal dynamics in job recommendation,” in *Proc. ACM Conf. on Recommender Systems*, 2026.
8. P. Nair, et al., “Acharya: A multimodal BERT & CNN and LLM agent architecture in job matching for ethical candidate,” in *Proc. IEEE Int. Conf. on Computing and Intelligence Systems*, Coimbatore, India, Oct. 2025.
9. T. Weber, et al., “Graph neural networks for candidate-job matching: An inductive learning approach,” *Data Science and Engineering*, 2025.



10. K. Tanaka, et al., “An intelligent resume–job matching system: Bridging candidates and opportunities through advanced semantic understanding,” in Proc. IEEE Int. Conf. on Intelligent Systems, Chennai, India, Jan. 2026.