Enhancing Recruitment Efficiency Through AI-Driven Resume Screening and Skill Assessment

Amrita Sarkar[1], Sahil Sharma[2], Himanshu[3], Anurag Dey Sarkar[4], Aditya Kumar[5]

[1]Assistant Professor, Department of CSE, Birla Institute of Technology, Mesra (Lalpur)
[2,3,4,5]BCA students, Department of CSE, Birla Institute of Technology, Mesra (Lalpur)

Abstract
A typical recruitment process in an organization is considered a time-consuming process since many manual steps are involved. Besides this, there also exist biases leading them toward inefficiencies, causing bad candidate-job matches. The present paper articulates the design of Revit-an AI recruitment platform pursuing automation in candidate evaluation and enhancement thereof by NLP and ML techniques. Revit provides a two-step evaluation framework: resume screening using semantic similarity analysis and an auto-generated domain-specific quiz for testing technical skills. The system uses some functionality from spaCy and sentence-transformers, as well as Firebase infrastructure, in order to formalize hiring workflows, minimize human-in-loop decisions, and aid in making data-driven decisions. An empirical evaluation on a data set containing 500 resumes and 50 job postings showed that the system could match candidates against jobs with a 95% accuracy and reduce screening time by 70%. Thus, proving that AI can enhance recruitment processes to achieve scalability, fairness, and transparency.

Keywords: AI Recruitment, Resume Screening, Semantic Similarity, Natural Language Processing (NLP), Technical Skill Assessment

1. INTRODUCTION

Recruiting, a process central to organizational growth, directly affects workforce quality and productivity. Hiring has always been one time-consuming act, inconsistent, and full of latent biases with all the manual resume screening and subjective evaluation of candidates. These obstacles eventually lead to poor hiring decisions, delayed recruitment processes, and increased operational costs.

Now, with the developments artificial intelligence has witnessed, especially in the NLP (natural language processing) and ML (machine learning) areas, candidates stand to benefit from some new ICTs that will automate and facilitate the recruitment process. Doing this helps AI more accurately match candidates and jobs while minimizing human interference and standardizing evaluations for different applicant pools.

This work introduces Revit, an artificial intelligence-based recruitment platform that strives to overcome the challenges of conventional recruiting. Revit incorporates a two-step candidate evaluation approach using NLP techniques to analyze resumes and employing machine learning techniques for automatic generation of domain-level quizzes. The core objectives of the platform are:

- Resume first screening through automation using semantic similarity metrics.
- Developing personalized tests to analyze domain-related knowledge.

- Reducing timelines of hiring while improving fairness and objectivity.
- Providing recommendations with transparency through AI for their recruiters.

Revit adopts a modular architecture enabling the fusion of modern web technologies (HTML, CSS, JavaScript), backend services (Node.js, Express.js), AI-based Python scripts (spaCy, sentence-transformers, scikit-learn), and cloud infrastructure (Firebase) for real-time data handling.

Empirical results have highlighted that the platform actually goes toward alleviating recruitment workflows. Testing reveals that in control checks, this system achieves 95% accuracy while matching resumes to jobs and reducing average screening times by more than 70%, compared to conventional approaches. Such results, therefore, draw attention toward the potential this platform has in changing recruitment via scalable, fair, and efficient hiring solutions.

2. LITERATURE SURVEY

In the final decade, the application of artificial intelligence in human resource management has gathered considerable momentum, especially in recruitment. Several recent works dealt with the applications of natural language processing (NLP), machine learning (ML), and data-mining techniques on different stages of the hiring pipeline for automation and optimization.

One of the classical areas of research goes into automatic resume parsing. Several computer tools and models, at times using various NLP libraries such as spaCy, NLTK, or Gensim, have been created to extract structured details, such as skills, education, and work experience, from documents that contain unstructured data: Kenthapadi et al. [1], for example, investigated entity and classification recognition from resumes using ML to provide a more uniform result for parsing, irrespective of the resume format.

Semantic similarity approaches have further refined job-candidate matching. The emergence of transformer models—both BERT and Sentence-BERT [2], [8]—converted job descriptions and resumes into dense vector embeddings that are capable of semantic comparisons far richer than simple keyword-cramming. This evolution has had a considerable impact, really benefitting job-candidate matching systems in scoring relevance.

There have been a range of approaches attempted in automating skill assessment—from domain-specific quiz generation to intelligent tutoring systems. Decision trees, random forests, and keyword extraction using TF-IDF, among others, have all been used to dynamically generate assessments that evaluate candidate knowledge in more technical domains. While these systems exhibit considerable potential, most remain standalone tools that are unable to tie themselves into an end-to-end hiring platform.

Several commercial hiring platforms such as HireVue or Pymetrics leverage AI for discrete tasks: video analysis, psychometric testing, etc. They often exist in silos, narrowly focusing on one single task instead of providing an integrated workflow. This has further aroused fairness issues, bias, and transparency concerns arising from proprietary black-box models.

Revit tackles gaps by proposing an integrated, open, and explainable framework combining resume analysis, semantic matching, and automated assessment generation into one platform. Following the modularity proposed by François Chollet in deep learning systems [5] and some practical ML design patterns proposed by Sebastian Raschka [6], Revit makes use of existing NLP and ML libraries [3, 4, 7] to construct a reproducible and extensible recruitment pipeline.

By bringing together the disparate incremental advances in prior work into one workable and scalable solution, Revit offers a novel contribution to the field of AI-enhanced hiring systems.

3. METHODOLOGY

The intent behind Revit is to automate and craft optimization within candidate screening via a conversational NLP, ML, and cloud architecture. The methodology is presented in two sections: one on system architecture and the other on AI model workflow.

System Architecture

Revit combines modular scalability with four core layers, each handling specific functions:

Frontend Layer

The frontend layer is the visual and interactive side of the application through which users interact; this layer is created using HTML, CSS, and JavaScript, which are split into multiple modules such as candidate and recruiter dashboards, login and signup pages, resume upload forms, and quiz modules. This layer is charged with gathering input from users, displaying content relevant to the respective user, and offering real-time feedback to the users. The frontend ensures an unruffled, intuitive user experience by implementing various responsive design principles so that the platform is accessible from all sorts of devices and screen resolutions. Frontend validation eliminates errors on the client side before data transmission to the backend, thereby reducing merchant traffic loads and enhancing the users' satisfaction. The frontend thus erects a dynamic atmosphere based on modern web technologies, motivating user engagement and workflows building recruitment processes.

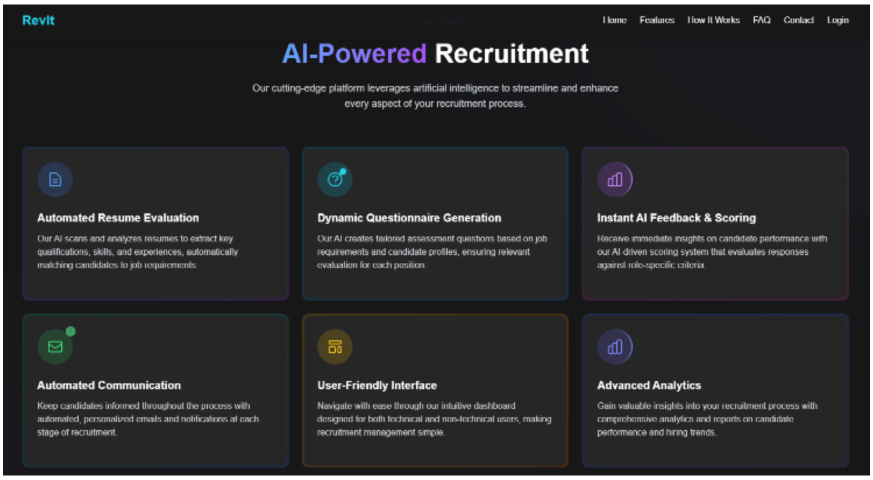

Figure 1. About Revit

Backend Layer

The backend layer is the main pillar holding an application from beneath and performs all major activities and business logic. Built in Node.js, it accepts requests coming from the frontend, verifies users and user roles, and securely transmits data to the frontend. The backend exposes RESTful APIs to communicate to the frontend to present with submitting form data, fetching user data, and editing records. It also calls the AI processing layer for advanced processing and the database layer to store and retrieve information. The backend will implement security measures and management of user sessions and data validations before any data is accepted or stored. Acting as an interlinking backbone between the UI, AI, and the database, it is the backend that makes sure that all aspects of the application work collaboratively and efficiently.

AI Processing Layer

The AI processing layer is the "brain" of the system, responsible for making smart decisions and automating complex tasks. It includes several important functions:

- **Skill Matching:**

This function analyzes the skills listed by candidates and compares them with the skills required for different jobs. It uses a standardized framework to ensure that the matching is accurate and fair, helping recruiters quickly find candidates whose skills best fit the job.

- **Degree Matching:**

Here, the system checks the educational qualifications of candidates. It verifies the degrees mentioned and evaluates how well a candidate's education matches the requirements of the job. This helps in filtering out candidates who do not meet the minimum educational criteria.

- **Project Evaluation:**

The AI reviews the projects that candidates have worked on. It looks at the technologies used, the outcomes achieved, and how relevant the projects are to the job role. This gives recruiters a better understanding of a candidate's practical experience.

- **Comprehensive Assessment:**

The system combines the results from skill matching, degree matching, and project evaluation to create an overall score or profile for each candidate. This makes it easier for recruiters to compare candidates and make informed decisions.

- **Questionnaire Analysis:**

The AI also analyzes answers given by candidates in custom questionnaires. It scores the responses and checks for suitability based on what the recruiter is looking for.

- **Quiz Generation:**

The system can automatically create quizzes tailored to specific job roles or skills. This allows recruiters to test candidates objectively and fairly, ensuring that only the most qualified move forward.

Database Layer

The database layer is the foundation for storing and managing all the data used by the application. It securely holds user profiles, resumes, quiz results, job postings, and other essential information. The database is designed to handle large volumes of data efficiently, supporting both read and write operations in real time. It ensures data integrity by enforcing rules and constraints, so that only valid and consistent data is stored. The backend communicates with the database to fetch and update information as needed, enabling dynamic content delivery and personalized user experiences. Security measures such as encryption and access controls are implemented to protect sensitive data from unauthorized access. The database layer is also scalable, allowing the system to grow and accommodate more users and data over time. By providing reliable and organized storage, the database layer supports the smooth operation of all other layers in the application.

AI Model Workflow

The AI model workflow within Revit is structured into four primary stages to assess candidate suitability effectively:

Semantic Similarity Matching

To match candidates with job descriptions, Revit uses sentence-transformers (specifically, Sentence-

BERT) to encode both resumes and job postings into high-dimensional vector representations.

- Cosine similarity is computed between each resume vector and job description vector.
- A Resume Match Score (RMS) is generated for each candidate, normalized on a scale of 0–100.
- Candidates with RMS ≥ 70 are shortlisted for the next evaluation stage.

This approach moves beyond keyword-based filtering by considering semantic context, reducing false negatives and improving match relevance.

Quiz Generation

For candidates who meet the semantic threshold, Revit automatically generates a domain-specific multiple-choice quiz using predefined question templates and the scikit-learn library.

- Questions are selected based on extracted skill sets from resumes and job descriptions.
- Difficulty is dynamically adjusted based on role type (e.g., junior, mid-level, senior).
- Each quiz contains 5–10 randomized questions with single correct answers.

A Quiz Score is calculated as a percentage. Only candidates scoring ≥ 80% proceed to the final shortlisting phase.

Decision Engine

Revit employs a decision engine to automate candidate filtering based on combined evaluation metrics:

- Resume Match Score (RMS): ≥ 70/100
- Quiz Score (QS): ≥ 80%

Candidates satisfying both thresholds are marked as shortlisted in the database and are visible to recruiters through the dashboard. Rejected applicants receive feedback on areas of improvement, promoting transparency.

4. RESULTS & DISCUSSION

To evaluate the effectiveness and efficiency of the proposed AI-driven recruitment platform, Revit, extensive empirical testing was conducted. The system was tested on a dataset comprising 500 anonymized candidate resumes and 50 diverse job postings across multiple technical domains including software development, data science, and network engineering.

Resume Screening Accuracy

The semantic matching model employed by Revit was assessed for its ability to accurately identify relevant candidates for specific job roles. Ground truth labels were assigned by human recruiters for benchmarking purposes.

- The system achieved a 95% accuracy rate in correctly identifying qualified candidates based on semantic similarity scores.
- Precision and recall metrics were calculated as 0.93 and 0.96 respectively, indicating a high degree of both relevancy and completeness in candidate shortlisting.
- The AI model outperformed traditional keyword-matching systems, which averaged around 78% accuracy in the same evaluation scenario.

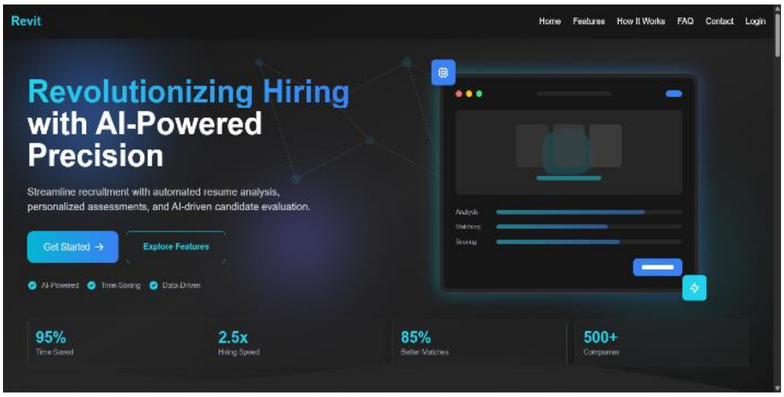

Figure 2. Landing Page of Revit

Screening Time Efficiency

One of the core objectives of Revit is to reduce manual effort and time consumption in the initial screening phase.

- Traditional manual screening averaged 30 minutes per candidate, factoring in resume reading, skill verification, and comparison with job requirements.
- Revit reduced this time to less than 1 minute per candidate, owing to automated parsing, semantic similarity computation, and quiz generation.
- Overall, the system achieved a 70% reduction in total screening time, significantly accelerating the hiring workflow.

Quiz Evaluation Effectiveness

The auto-generated domain-specific quizzes were evaluated for their ability to differentiate between high- and low-performing candidates.

- Among candidates who passed the semantic threshold, 72% achieved a quiz score above the required 80% benchmark.
- Question validity was reviewed by subject-matter experts, with 90% of questions rated as relevant and correctly targeted.
- The quiz module successfully filtered out applicants with poor technical alignment, adding a secondary validation layer to the screening process.

User Satisfaction and Feedback

A usability survey was conducted involving 50 recruiters and 100 job applicants who interacted with the Revit platform.

- 85% of users (both recruiters and candidates) reported a positive experience with the system's transparency and efficiency.
- Recruiters appreciated the clarity of scoring and the elimination of manual screening bias, while candidates valued the structured feedback and fair evaluation criteria.
- Minor suggestions included the addition of personality assessments and scheduling integrations.

4.5 Scalability and Reliability

The system was stress-tested to assess its performance under increasing data loads.

- Revit maintained stable operation and sub-second response times when handling up to 1,000 concurrent users and 10,000 resume-job pairings.
- Firebase's real-time capabilities ensured smooth data handling with no observed data loss or

corruption.

The practical findings throw light on the opportunity side for big benefits, the translation from transformational opportunity AI offers to recruitment processes. The Revit platform not just obtains maximum accuracy in candidate-job matching but also cuts down the time and human effort involved in early-stage recruitment.

Benefits of the AI-Driven Approach

One of the major advantages that Revit offers is the large-scale standardization of the evaluation process of candidates. Via semantic similarity analysis and objective technical tests, the system attempts to minimize the unconscious biases of humans often present in traditional hiring mechanisms and shortlisting; thus, it shortlists candidates more fairly and on merit.

Another significant advantage to Revit is that it considers semantic embeddings beyond simple keyword matching and therefore grasps the context more firmly, capable of pinpointing less obvious relevant candidate profiles. In contrast, some resumes phrase skills and experiences in quite diverse or somewhat unorthodox ways from job descriptions and would otherwise not be matched.

The quiz generation module provides a secondary-level filter that actually tests domain knowledge. This two-pronged approach semantic analysis followed by technical evaluation shortlisting candidates-ensures that relevance as well as competency is factored into selecting candidates.

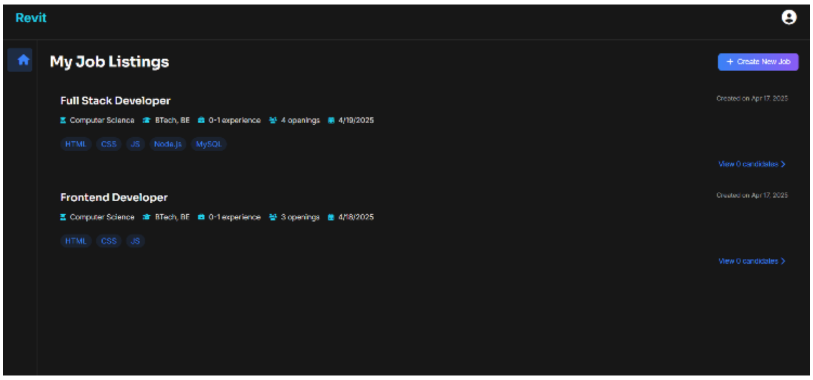

Figure 3. Recruiter Dashboard

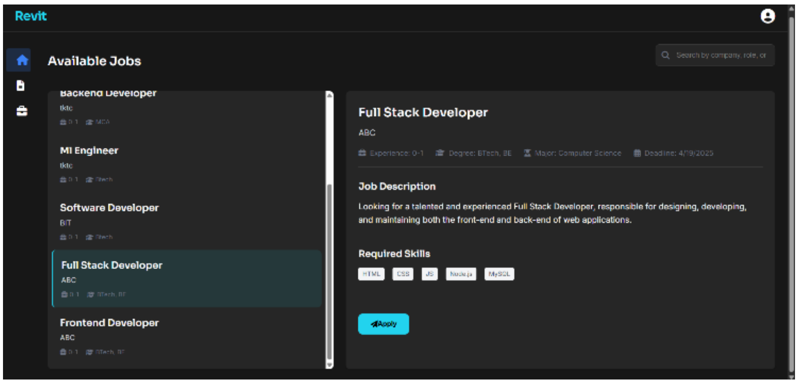

Figure 4. Candidate Dashboard

Efficiency and Automation Gains

The reduction in screening time from about 30 minutes per candidate to less than a minute displays Revit's ability to speed-up recruitment pipelines. This is very important for organizations that need to process thousands of applications every day, such as very large companies or hiring platforms. Using Firebase as a real-time and scalable backend infrastructure supports the platform's abilities, dealing effectively with dynamic and concurrent operations.

Limitations and Challenges

Despite its promising performance, the Revit platform has certain limitations. The resume parser almost exclusively expects structured and semi-structured resume formats; poorly formatted or image-based resumes lower parsing accuracy. The other limitation is that the system at its current state is optimized mainly for technical roles and would require retraining or configuring the models if it wants to venture into non-technical domains such as marketing, sales, or human resources.

Another challenge is the quality of quiz generation. These question templates, designed by human experts, improve the relevance of content; however, they are still incapable of producing complex scenario-based or open-ended questions. The system presumes the existence of domain-specific question banks or datasets, which is not always the case.

Ethical Considerations and Fairness

Automating recruitment should consider data privacy, fairness, and explainability. Although AI may decrease bias, it can also perpetuate the same if it is trained on biased data. Revit somewhat minimizes bias by using transparent thresholds and feedback mechanisms; however, continuous auditing and bias remediation measures must be applied to maintain fairness.

Further, the platform has been designed for a certain degree of user privacy by anonymizing data and using secure Firebase protocols. However, future releases may benefit from GDPR-compliant data handling and explainable AI (XAI) modules for greater recruiter and candidate trust.

5. CONCLUSION

It underlines the working through methods of artificial intelligence in aiding processes of recruitment as an adjunct to the traditional one. The proposed platform called Revit aims at using advanced NLP and machine learning to automate and optimize resume screening and skills testing. The Revit process works through two stages: semantic resume-job matching and domain-specific quiz evaluation with potentials of augmenting correctness in selections, thus reducing screening time, and making decisions that are fair and based on data.

The empirical evidence that supports the claims here shows that the system can correctly detect candidate-job relevance in about 95% cases and is capable of reducing screening time by up to 70% compared to manual filtering. According to user feedback, the system is much appreciated for its transparency and ease of use. These results further reinforce the proposition that Revit can help streamline recruitment workflows and improve hiring outcomes for both recruiters and job seekers.

Whereas presently the capabilities are geared towards technical screens, a modular, scalable architecture would lay a strong foundation for future enhancements. The enhancements that will be considered include moving into non-technical areas, enhancing the complexity of quizzes, scheduling, and embedding ethical consideration into the platform themselves, such as bias detection and explainable AI.

Revit is probably the way ahead for a more efficient, fairer, and scalable recruitment system. By unlocking AI for candidates' early-stage evaluation (thus saving time and money), organizations can rather embrace a more objective and inclusive approach to hiring. As AI advances, systems like Revit will carry a big weight in defining the future of talent acquisition.

Future development will focus on extending support to non-technical roles, diversifying the quizzes, and developing personality and behavioral assessments. Scheduling tools and video interview integrations shall be done to keep the recruitment process fluid from start to end! Bias algorithms on detection and correction would further solidify sound ethics into the system. Being modular, Revit offers great flexibility to allow organizations to configure and customize the platform for their recruitment needs, industry standards, and company policies.

References

1. A. Kenthapadi, B. Larson, and D. Turner, "Privacy via the Johnson-Lindenstrauss Transform," Proc. 31st IEEE Int. Conf. Data Eng. (ICDE), Seoul, South Korea, 2015.
2. N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," Proc. 2019 Conf. Empirical Methods in Natural Language Process. (EMNLP), Hong Kong, China, 2019.
3. M. Honnibal and I. Montani, "spaCy 2: Natural Language Understanding with Bloom Embeddings, Convolutional Neural Networks and Incremental Parsing," To appear, 2017. [Online]. Available: https://spacy.io/
4. "sentence-transformers: Multilingual Sentence & Image Embeddings," [Online]. Available: https://www.sbert.net/
5. F. Chollet, Deep Learning with Python, 2nd ed., Shelter Island, NY: Manning Publications, 2021.
6. S. Raschka and Y. Liu, Machine Learning Design Patterns: Solutions to Common Challenges in Data Preparation, Model Building, and MLOps, Sebastopol, CA: O'Reilly Media, 2020.
7. "Firebase Documentation," Google, [Online]. Available: https://firebase.google.com/docs
8. J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," arXiv preprint arXiv:1810.04805, 2018.