IARJSET
ISSN (O) 2393-8021, ISSN (P) 2394-1588
International Advanced Research Journal in Science, Engineering and Technology
Impact Factor 8.066 ⁒ Peer-reviewed & Refereed journal ⁒ Vol. 13, Issue 7, July 2026
DOI: 10.17148/IARJSET.2026.13743

AI-Based Resume Screening and Job Matching Systems: A Comprehensive Review of NLP, Machine Learning, Explainable AI, and Fairness-Aware Recruitment Approaches

Naveen J[1], Manoj M[2], Jayanth Vishwakarma G[3]

Department of MCA, Surana College (Autonomous), Bengaluru, India[1-3]

Abstract: Now machines help find workers fast - no need to shuffle piles of paper anymore. Instead of people flipping through endless applications, software reads them in seconds. Because it works quicker, teams avoid delays that slow down hiring. With fewer eyes on each detail, some errors slip away quietly. Machines compare abilities to job needs without tiring. This shift lets staff spend energy talking to promising applicants instead of sorting names. Tasks once done by hand now happen before lunch. Focus moves where it matters most: real conversations after filters do the math. Now machines grasp what words really mean inside a resume, moving past just spotting keywords. Because of smarter algorithms, connections among skills, jobs held, degrees earned, licenses obtained, and role requirements become clear. These links let software spot people who fit well without needing every phrase copied exactly. Starting with patterns found in massive amounts of text, systems build understanding that matches human judgement more closely. Hidden layers within networks detect subtle clues others miss during screening. As training improves, so does precision when sorting applicants by relevance. Less time gets spent reviewing unfit profiles thanks to these behind-the-scenes shifts. Even with progress, hurdles still block broad use of AI in hiring. Problems like skewed algorithms pop up when old hiring records shape new tools. Decisions made by machines often hide how they work, raising red flags around fairness. Personal details of job seekers can slip into grey zones without clear safeguards. Outcomes might tilt against certain groups if past imbalances feed the system. Making sense of AI choices matters more now than ever before. Efforts focus on building methods that reveal their reasoning step by step. Fairness built into code helps balance the scales where bias once lived unseen. Trust grows only when processes open up to inspection. Researchers dig into ways that show not just results, but why those results happen. This piece looks into recent studies about how artificial intelligence helps sort resumes and match people to jobs, using tools like TF-IDF, cosine similarity, transformers, and small neural setups, along with big language models. Each method shows distinct pros and cons when tested in varied hiring environments – some work faster, others understand context better. Efficiency gets a boost through AI, alongside sharper matches between applicants and roles, raising the standard of who gets hired. Still, what comes next needs clearer logic behind decisions, room for human judgment, strong ethics, plus ongoing checks for unfair patterns. When smart tech teams up with mindful choices, companies might shape hiring systems that perform well while staying just, open, and dependable - for those offering jobs as much as those seeking them.

Keywords: Artificial Intelligence, Resume Screening, Job Matching, Natural Language Processing, Machine Learning, Explainable Artificial Intelligence, Recruitment Automation, Fairness, Semantic Similarity, Large Language Models.

I. INTRODUCTION

Hiring people sits at the core of what companies do, given how worker skills shape creativity, output, and staying power over years. With firms reaching across borders and more jobs posted online, each opening now pulls in stacks of applications from everywhere. Even though applying takes just clicks today, sorting through names and backgrounds floods hiring teams with work they cannot handle alone. Reading every resume by hand drains energy, slows progress, introduces errors, reflects hidden preferences, and lifts expenses higher. So tools driven by smart systems enter the scene - not to replace judgment but to lighten loads without losing care or balance. Out of nowhere, machines began reading resumes like humans do. Instead of just scanning for keywords, they now grasp what words mean together - thanks to tricks pulled from natural language processing. Picture this: a system learns patterns in job histories, picks up skills without being told, then lines them up against open roles.

As it turns out, sorting applicants doesn't need rigid rules anymore. Hidden connections between education and experience get uncovered through layers of learning models. Recruiters see better matches come forward - not because something was programmed step by step - but because the machine figured it out over time. What once took hours now

DOI: 10.17148/IARJSET.2026.13743

happens in moments, cutting down busywork dramatically. Hiring isn't about ticking boxes now; it's guided by insights drawn from messy, real-world text. Slowly but surely, old tracking tools are fading into background noise Out of nowhere, smarter hiring tools started emerging thanks to new kinds of language models. Instead of just matching words, BERT and its variations began grasping meaning through context, unlike older methods stuck counting terms. Because they understand sentences more like people do, these systems spot connections between job ads and applicants even if phrasing differs. Picture an engine that links a person's background to open roles by sense - not keyword repeats.

Over time, matches got sharper, helping companies find better fits without needing exact phrases. This shift didn't just boost precision - job searches now feel less generic, smoother for everyone involved Even with today's tech progress, using artificial intelligence in hiring brings up real concerns about fairness and usability. Because machine learning studies past hiring choices, it can pick up on unfair patterns tied to gender, race, income level, or schooling. When those skewed patterns live inside the data used to train AI, the system might repeat biased outcomes without meaning to, deepening gaps already present. What adds to the problem is that complex neural networks often work like hidden engines - recruiters and job seekers alike struggle to see why certain decisions come out the way they do. Now worries grow, scientists look into clear AI, fairer algorithms, plus hiring tools that show their reasoning without losing accuracy. Because of this shift, making smart recruiting both efficient and open stands out as a key challenge today When things started shifting, one look at what papers say shows how AI now handles resumes and matches jobs. Not only do computers read words, they find patterns using methods like NLP or deep learning networks that mimic brain activity. Some tools even explain why decisions happen, bringing transparency into hiring steps. Instead of backing just one method, this work pulls together many results to show where the field really stands. Hidden flaws pop up often - bias is common - even when systems claim neutrality.

Old ways get compared here, but attention turns quickly toward what gets missed in studies so far. Ideas change fast; today's standard might fade by next year. What matters most? Real impact on companies trying to hire right without repeating past mistakes This review aims to do four main things. One path it takes is looking into how artificial intelligence has changed when picking resumes and linking people to jobs. Instead of just listing pros, it digs into what current smart hiring tools can handle well - yet also where they fall short. Another direction involves checking whether these automatic systems treat everyone fairly, make their choices clear, show how decisions happen, and stick to strong moral standards. Looking ahead, fresh paths emerge - ways that could shape fairer, clearer hiring tools shaped around people. Instead of scattered details, a full picture forms here: one pulling together new breakthroughs in automated hiring methods. This clarity helps those building, studying, or using these systems grasp both what works and where risks hide. With balance in mind, progress gains direction - not just speed. Trust grows when choices are open, grounded, guided.

II. EVOLUTION

Thirty years ago, hiring looked nothing like it does today. Computers started changing everything slowly at first, then faster than anyone expected. People used to sort through paper CVs by hand, one after another. Judgment came from experience, yet the process dragged on for weeks. Costs piled up without warning. Scaling meant adding more staff, which only magnified delays. Global companies faced a flood of interest whenever they posted roles. Online job boards exploded in number. A single position might pull in five hundred replies overnight. Reviewing each applicant felt impossible under tight deadlines. Errors crept in when fatigue took hold. Efficiency became harder to maintain across continents. Machines stepped into the gap where humans could not keep pace. Software learned to filter keywords before eyes ever saw a resume. Automation wasn't welcomed everywhere at once. Some doubted machines could grasp subtle strengths. Others feared losing personal touch. Yet pressure mounted regardless. Systems evolved because necessity demanded solutions. What emerged reshaped how talent finds work now. One big step forward in hiring tech came with tools called Applicant Tracking Systems. Resumes started living online thanks to these platforms, letting companies sort applicants using set keywords instead of paper files. Sorting through candidates sped up a lot since unqualified profiles got weeded out quickly. Admin work dropped sharply once machines handled first-round checks alone. Exact word matches ruled how those older versions worked - little room existed for variation back then. So resumes got tossed out a lot when they didn't match exact keywords, even if skills were there just phrased another way. That flaw made hiring tools less useful, pushing experts to hunt for smarter methods that actually grasp what words mean in context Something shifted when machines started grasping how people really write. Instead of just scanning for exact words, they began making sense of meaning within sentences found on CVs and job posts. A computer now pulls out names, schools attended, abilities listed, credentials held, past jobs done, phone numbers - all without humans touching each page. What helped most was understanding that similar ideas can be phrased differently across texts. Matching got smarter because connections between roles and talents were seen beyond repeated terms alone. Picking applicants turned more precise, fewer strong fits slipped through unnoticed When hiring data grew larger, companies started using machine learning more often to make better decisions about candidates. Instead of sticking to fixed rules set by people, these smart

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Impact Factor 8.066 ⨳ Peer-reviewed & Refereed journal ⨳ Vol. 13, Issue 7, July 2026

DOI: 10.17148/IARJSET.2026.13743

programs found trends in past hiring records to judge how well someone might fit a role. Tools like Random Forest, Support Vector Machine, Logistic Regression, Naïve Bayes, and K Nearest Neighbours showed strong results sorting applicants, rating resumes, and recommending hires. Thanks to machine learning, businesses built custom suggestion systems that matched job seekers with openings suited to their backgrounds, making the whole process faster and smoother for everyone involved. Deep learning changed hiring tools by helping machines grasp meaning in written text. Instead of relying on old number-crunching methods, neural nets like CNNs, LSTMs, and transformers began reading context far better. BERT showed up later - then SBERT followed - both standing out when turning resumes and job posts into smart word patterns. When words shift across documents, these models still link similar ideas through surrounding phrases. Because they work so well, many current AI recruiting platforms now run on them. Lately, big language machines changed how hiring gets automated. These tools understand words better than older ones made just for single jobs. Because of that power, they summarize resumes smartly, build worker profiles, help during interviews, suggest careers, even chat through hiring steps. Reading tough job posts comes naturally now. They spot missing skills, draw out hidden patterns about people applying, give clear reasons behind choices humans make when picking hires. Still, more use brings bigger worries - about bias, secrecy, high computing demands, whether results can be trusted. So scientists today lean toward mixing transparent AI methods, fair-training tricks, plus ways to keep people involved while tech helps decide who moves forward. One step at a time, hiring tools powered by artificial intelligence have moved from human-driven choices toward smart systems that understand context and can justify their decisions. Today's platforms mix machine learning with natural language processing, deep learning alongside AI methods designed to show how conclusions are reached - automating assessments without ignoring concerns about bias or clarity. Change didn't arrive overnight; instead, it reshaped how organizations find talent, setting patterns likely to influence hiring far ahead.

III. LITERATURE REVIEW

Lately, Artificial Intelligence grabbed the spotlight in hiring studies thanks to its ability to handle resume sorting automatically. Instead of people reading every application, machines began stepping in to lighten the load for recruiters. At first, these tools leaned on simple keyword checks inside Applicant Tracking Systems. Speed improved, true, yet meaning often slipped through the cracks - similar skills missed if phrased differently. Contextual gaps like those pushed experts toward smarter methods. Slowly, focus shifted to Natural Language Processing and Machine Learning, aiming to grasp not just words but what they simply.Understanding intent became key, letting systems compare candidates and roles beyond mere phrases. Smarter parsing emerged, shaped by how humans actually describe experience and needs. Efficiency evolved, not just from speed but from sharper judgment. Ideas once limited to matching terms now reach into meaning itself. Progress unfolded quietly, driven less by flash than function. Today's approaches reflect that quiet shift - one line of reasoning at a time. Long before smart hiring tools became common, systems began pulling data from resumes automatically through language-processing methods. Instead of reading each document by hand, computers learned to spot degrees, abilities, past jobs, licenses, and career wins hidden in messy formats. Muthu Kumar built a tool like this, mixing language analysis with learning models to judge applicants more consistently. It matched job summaries against CVs, giving them score rankings based on how closely they aligned. To sort better, the design used word frequency patterns paired with angle-based comparisons across texts. Even though people saw time saved, results still relied on surface-level wording counts missing richer meanings behind skill relevance. One way past basic keyword searches came through new methods built around meaning instead of word matches. Using artificial intelligence, Ajjam and Al-Raweshiday designed a system where resume content met job needs via TF-IDF scores and cosine math to spot similarities. Experiments showed this idea beat older techniques because it linked ideas even when phrasing differed. Matching meaning helped hiring see deeper into fit - especially useful in tech roles where people describe the same skill many ways.

Less chance of losing good applicants just due to wording gaps was one clear benefit noticed Fast progress in transformer language systems gave hiring tools a boost. A team led by Akhtar built an AI tool for suggesting resumes, using BERT to understand meaning in text. Instead of old number-focused approaches, these new models grasp context from both directions in language, helping match applicants to roles better. In much the same way, newer research uses Sentence-BERT and similar deep learning designs to sharpen how closely texts are related in sense. They get stronger results than earlier natural language methods, especially when jobs demand mixed skills or varied expertise is listed across different fields. Better hiring tools aren't just about precision - fairness matters more now than before.

Old data often carries hidden patterns tied to gender, race, age, or school history. When algorithms learn from that, they might repeat old unfair choices without meaning to. One path forward involves making machine decisions easier to follow - peering under the hood, so to speak. SHAP helps show which details sway a hire, spotlighting what drives each outcome. Alsubaie and Aleisa suggested this transparency could reshape how systems pick applicants. One study showed clear explanations help hiring teams feel more confident while giving coders better tools to spot unfair patterns in automated

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Impact Factor 8.066 ⦂ Peer-reviewed & Refereed journal ⦂ Vol. 13, Issue 7, July 2026

DOI: 10.17148/IARJSET.2026.13743

choices. Other research looked at methods like LIME, learning models tuned for equity, and open systems designed to weigh accuracy against fair hiring - without tipping too far either way. Even with clear advances, some problems still show up in the research. Many current papers judge hiring tools by how accurate they are but skip over issues like bias, responsibility, clarity, or real-world use. Without a shared standard dataset, comparing one system to another stays tricky. New steps forward with large language models open doors for smarter ways to review applicants, though heavy computing needs, data safety risks, unclear reasoning paths, and shaky oversight rules slow things down in business settings. Work ahead must aim at building selection methods that stay open, balanced, involve people meaningfully, deliver strong results without cutting corners ethically.

IV. COMPARATIVE ANALYSIS OF EXISTING STUDIES

Author / Study	Primary Technique	Major Contribution	Advantages	Limitations
Albaroudi et al. (2024)	Fairness-aware AI	Bias mitigation in recruitment	Improves ethical hiring	Limited practical implementation
Akhtar et al. (2025)	BERT + Recommendation Engine	Context-aware candidate recommendation	High semantic accuracy	High computational requirements
Muthu Kumar (2025)	NLP + TF-IDF	Automated resume analysis	Simple and efficient	Weak contextual understanding
Ajjam & Al-Raweshidy (2026)	Semantic Similarity	Context-aware job matching	Better recommendation quality	Dataset dependency
Desai et al. (2025)	NLP + React + AI	Intelligent resume screening	Fast candidate ranking	Limited fairness analysis
Yadav et al. (2025)	NLP + ML	Career guidance and recruitment optimization	Improves recruiter productivity	Limited explainability
Alsubaie & Aleisa (2025)	Explainable AI	SHAP-based bias detection	Transparent recruitment	Additional computational cost
Dogra et al. (2025)	NLP + ML	Automated resume ranking	High screening accuracy	Requires quality datasets
Anusha et al. (2025)	ML + NLP	Resume classification	Improved recruitment automation	Traditional ML limitations

Table. 1 Comparative Analysis of Existing AI-Based Recruitment Studies

V. TECHNOLOGIES USED IN AI RECRUITMENT SYSTEMS

Starting off, AI-powered hiring tools bring together different computer methods to handle tasks like reviewing resumes and pairing people with jobs. Instead of depending only on fixed rules or simple word searches, these newer systems weave in advanced techniques such as artificial intelligence, natural language processing, machine learning, deep learning, transformer models, and explainable AI. Together, they help sort applicants more precisely while cutting down hands-on work. One key benefit is better alignment between a role and who fills it - thanks to smarter analysis behind the scenes. Though each method plays its own part, their joint effect builds something stronger: a smarter way to hire. What sets them apart isn’t just speed - it’s how clearly they show why one person fits over another.

A. ARTIFICIAL INTELLIGENCE(AI):

Something smart machines do is handle parts of hiring work once done only by people. Because computers learn patterns, they sort through job applications quickly. Resume checks, picking top applicants, setting up interviews - these steps now happen faster. One reason companies adopt these tools lies in how evenly choices get made across hundreds of profiles. When software predicts who might succeed, managers save hours usually spent reading every single document. Speed meets fairness when tech takes on early-stage evaluations. Tools driven by learning algorithms change how teams find new hires without slowing down. Learning never stops for AI when it comes to hiring data, spotting trends tied to good hires over time. Because they review degrees, past jobs, credentials, skill sets, along with work history, these tools point toward candidates matching what a company truly needs. With routine chores handled automatically, attention shifts where it matters most - planning next steps instead of paperwork. So now, intelligence built into software quietly shapes how HR operates across today’s tech driven job markets.

DOI: 10.17148/IARJSET.2026.13743

B. NATURAL LANGUAGE PROCESSING (NLP):

Most hiring tools today rely on Natural Language Processing since job details and applicant files come as messy text. Because computers struggle with free-form writing, they need help pulling meaning from words strung together without clear rules. Through NLP, machines break down sentences into organised data points they can actually work with. What looks like ordinary paragraphs to us turns into patterns a system can compare, sort, and learn from. Without this step, automated screening would not handle real-world applications effectively. Inside hiring tools, natural language processing handles key jobs - think pulling data from resumes, breaking text into parts, reducing words to roots, finding real-world items like job titles, spotting abilities listed, uncovering past roles, checking degrees held, plus grasping meaning behind phrases. Picking apart CVs pulls out organized facts: names of applicants, phone numbers, schools attended, credentials earned, tech know-how shown, coding tongues known, employment timelines - all lifted straight from PDFs or Word files. That gathered info then feeds into lining up people with open positions they might fit. Because of today's NLP methods, machines grasp meaning by spotting how words connect in context. Take job titles like "Software Engineer," "Application Developer," or "Software Programmer" - systems now see these refer to similar jobs despite varied terms. Unlike old-school tools that just match keywords, these approaches get nuance. Better matches happen when meaning matters more than exact phrasing.

C. MACHINE LEARNING (ML):

Learning by machines helps hiring tools pick up patterns from past choices rather than just following fixed instructions. Because it studies what traits linked to good hires before, the system guesses who might work out well now. Patterns found in old data shape how new applicants get assessed over time. Instead of rigid checklists, software adapts using examples drawn from earlier decisions. What worked in the past quietly guides future picks through repeated exposure. The closer someone resembles prior successes, the more likely they appear suitable. This way, predictions grow sharper without constant human tweaking along the way. Among tools in hiring studies, supervised methods show up often. Though less flashy, Random Forest stands out due to steady results and strong performance spotting applicant patterns. When dealing with dense resume texts, Support Vector Machines handle separation tasks without much fuss. Logistic Regression gives clear yes-or-no signals, making it useful for straightforward go/no-go choices. Simpler models like K-Nearest Neighbours stay around - not fast talkers but quick workers - thanks to low setup needs. Naïve Bayes fits the scene too, offering speed without heavy tuning.

D. DEEP LEARNING(DL):

One layer builds on another inside these systems, letting patterns emerge without handcrafted rules. Where older methods depend on human-designed traits, this approach pulls structure straight from piles of paperwork by itself. Hidden details rise through stacked processing steps, making sense in ways simpler models miss. Accuracy grows because meaning gets captured across levels, not just at surface points Starting off, various neural network designs show strong results in hiring-related studies. Instead of just linking words, CNNs pull out fine-grained details from resume text. Moving forward, LSTMs handle order-based patterns, capturing how ideas connect across a document. Surprisingly, transformers now lead the field - not due to trendiness but stronger grasp of meaning and context in job applicant data Though deep learning boosts how resumes get sorted, picks job seekers, judges meaning matches, reviews interviews by machine, guides careers smarter - big computing power is usually needed. Training takes heaps of data, runs slow, asks more time than standard methods ever did.

E. TRANSFORMER MODELS: BERT AND SENTENCE-BERT:

Words before and after get equal attention when BERT reads a sentence. This shift changed how machines grasp what we mean. Earlier systems moved left to right, missing connections across phrases. Because it sees full context at once, meanings become clearer. Breakthroughs like this don't happen often in language tech Resumes and job posts get turned into smart word patterns by BERT when hiring. Because of this shift, people with similar abilities - though described differently - are linked more precisely. Matching isn't stuck on exact terms anymore; meaning matters now instead. Better fits come through quietly, without shouting about keywords Fixed-length sentence embeddings come from Sentence-BERT, an upgrade to BERT built for spotting meaning matches. When picking resumes, matching applicants to roles matters a lot - this is where SBERT shows up often in hiring tech. Instead of just counting words, it grasps context better than older TF-IDF approaches. Tests keep showing BERT variants rank candidates more accurately. Performance gains appear across understanding intent, sorting fits, and suggesting hires Even with their benefits, transformers need strong computing power along with large memory space. Because of this, teams keep exploring simpler versions that can run faster in factory settings.

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Impact Factor 8.066 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 13, Issue 7, July 2026

DOI: 10.17148/IARJSET.2026.13743

F. EXPLAINABLE ARTIFICIAL INTELLIGENCE (XAI):

Even though computers now help sort job applicants faster, some smart programs work like hidden boxes - nobody sees how they pick people. Without clear answers, those left out are stuck wondering why they did not get chosen Because artificial intelligence can be hard to understand, Explainable AI gives clear reasons behind its choices. Recruiters see what drives each hiring suggestion when results come through. What makes a candidate stand out becomes visible thanks to these tools. One method, known as SHAP, breaks down how every detail affects the outcome. Another approach, called LIME, builds simpler models around individual cases. These ways help people trust automated systems just a little more One way to look at it: SHAP measures how much each applicant feature affects hiring choices. Meanwhile, near specific cases, LIME builds basic models to show why a decision was made for one person. Recruiters feel more sure about outcomes when they see these explanations. Rules are easier to follow because the process becomes clearer. Checking for fair treatment gets simpler too. Developers gain insight into unnoticed patterns linked to demographics hidden in past data. When machines help pick job candidates, clear explanations matter. Not magic boxes deciding alone - people stay involved when they see how answers form. Seeing the why behind choices builds confidence slowly over time. Fairness grows stronger if logic shows up plainly. Trust rises once hidden steps turn visible. Systems work better alongside humans who question, adjust, learn. Understanding replaces guessing. Decisions open up without drama. Clarity becomes normal. Judgement shifts - but does not vanish.

VI. GENERAL ARCHITECTURE OF AI-POWERED RESUME SCREENING SYSTEM

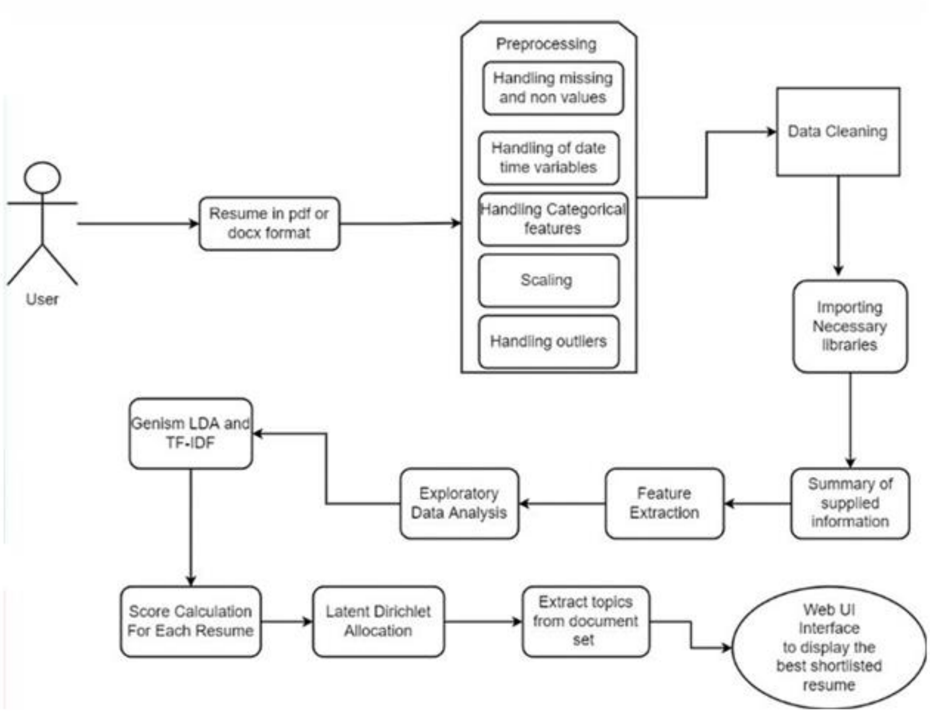

Fig. 2: General Architecture Of AI-powered Resume Screening System

VII. FAIRNESS, BIAS, AND ETHICAL CHALLENGES IN AI RECRUITMENT

When machines help sort job applications, they speed up who gets picked first. Yet depending too much on smart software brings new problems around what is fair, lawful, or right in how people are chosen. Even though computer decisions seem neutral, they copy patterns from old hiring records - records that sometimes favor certain groups. Because of this, hidden preferences from the past might carry forward without anyone noticing. Making sure these tools stay clear, honest, responsible, and respectful of personal information now matters more than ever. Figuring out how to build trust inside automated hiring shapes much of today's work in this field.

A. Algorithmic Bias

Most of the time, skewed outcomes in automated choices come from how machines were taught to decide. Picture a system that picks job candidates - it might follow old habits pulled straight from past company records. When earlier

DOI: 10.17148/IARJSET.2026.13743

hires leaned heavily toward certain genders or ethnicities, those preferences stick around inside the software. Past favoritism linked to age, school history, or income level can quietly shape what the model suggests next. Learned behavior shows up again when new applicants get reviewed under hidden assumptions. Machines do not invent prejudice - they copy it, silently, from what they were fed before Take old hiring files filled mostly with men in coding jobs. A machine learning tool might then link tech skills more closely to men, quietly rating women with equal experience lower. Start elsewhere - say, a model shaped by resumes from just a few colleges. It could lean toward graduates from those schools, passing over skilled people from smaller or lesser-known programs When algorithms favor certain candidates unfairly, hiring feels less fair. This unfairness can lead to lawsuits, public criticism, or teams that lack variety. Because of these outcomes, spotting skewed patterns now drives much of today's work in AI-based hiring tools.

B. Fairness in AI-Based Recruitment

Bias-free hiring means machines judge people by their abilities, not who they are. When a system looks at training, past jobs, or hard skills instead of age, gender, or background, it stays neutral. Decisions rooted in actual work fit keep things level. Judging someone by what they can do matters more than where they come from. Performance links to role demands, never identity markers. Equal treatment begins when software ignores traits that have nothing to do with doing the task Now imagine checking if hiring tools treat everyone fairly. Some ways to look include how often each group gets picked, whether chances are similar for qualified people, if outcomes stay balanced across groups, if predictions hold up equally, and whether one group faces steeper hurdles. Each of these checks helps spot patterns where an AI might favor certain applicants over others. Watching these signals shows whether results shift depending on who is applying One way to make hiring fairer? Adjust how machines learn, shifting either the data they train on, their decision rules, or final choices - aimed at cutting bias. Some methods reshape dataset balance, others pit models against biased patterns, add fairness limits during learning, or reframe traits so differences based on background fade without hurting performance Even when models aim to be fair, getting the balance right with accuracy is still tricky. Pushing too hard on fairness can weaken how well predictions work. This tension keeps researchers busy looking for better ways.

C. Transparency and Explainability:

Most people want to know how hiring decisions are made. Some modern systems learn patterns so intricate that even experts struggle to explain their choices. Because these tools hide how they work, job seekers often get no clear reason when passed over. Explaining rejections becomes a challenge when the machine's reasoning stays hidden Human-friendly insights emerge when AI shows how it picks job candidates. One method measures what parts of a resume weigh more in its choice, breaking down each factor's role step by step. Another approach zooms into single decisions, sketching clear summaries anyone can follow. Each technique reveals hidden logic without needing expert knowledge What shows up inside the model helps recruiters trust choices more when they see what traits sway picks strongest. Seeing how it works lets builders spot quiet unfairness plus check if hires lean on real job fit instead of age, gender, or background With shifting rules on artificial intelligence, clear explanations matter more. Because trust depends on showing how systems work. So openness helps prove that decisions are fair. When guidelines change, being upfront builds credibility. Since accountability can't be an afterthought. Thus visibility into processes becomes essential. If organizations want to show they're acting responsibly.

D. Data Privacy and Security:

Personal data like names, job history, pay goals, and ID papers move through hiring platforms every day. Because these facts are so private, handling them needs serious care. Information flows in many directions - gathered here, saved there, used somewhere else - all while risks grow quietly beneath the surface. Mistakes or weak safeguards can expose lives instead of protecting them. How firms manage access, where they store files, who sees what - each step shapes safety more than any single policy ever could Secret job application details could leak if hackers break into hiring systems, opening doors to stolen identities or misused private records. Because of this risk, artificial intelligence tools used in hiring need tough digital shields - like scrambled data storage, tight login checks, limited user permissions, hidden personal identifiers, and frequent safety reviews - to stay safe. What matters most is protection that works without drawing attention to itself Compliance isn't optional - global rules like GDPR set strict boundaries. Think India's DPDP law, shaping how personal data gets gathered. Rules differ by country, yet ethics remain central. Handling info? Laws expect fairness, transparency. Each nation adds its layer, making adherence complex. Still, ignoring them isn't an option. Privacy demands respect, no matter the region Fed by growing concerns over data leaks, methods like federated learning along with differential privacy now offer a way to shield applicant details

DOI: 10.17148/IARJSET.2026.13743

without dulling how well models work. Though fresh on the scene, these approaches quietly adjust training so sensitive inputs stay hidden across devices instead of being collected centrally. Their rise ties closely to stricter rules and rising distrust in bulk data handling, pushing researchers toward setups where insight grows minus exposure.

E. Ethical Considerations:

Built into smart hiring tools, ethics matter - lives shift when jobs are gained or lost. Fairness does not appear by accident; it grows where systems choose transparency over speed. One choice at a time, inclusion takes root when bias is questioned, not ignored. Real names behind applications deserve dignity, not automated guesses. Decisions ripple outward, so checks and awareness must be part of each step forward What stands out most is how heavy dependence on machines can blur moral lines when people step back too far. Even though algorithms lend a hand in sorting choices, real judgment belongs with seasoned professionals who see what data cannot. Only humans catch the subtle details that code might miss. Important calls need eyes that notice more than patterns. Decisions shaped only by logic risk ignoring life's messy edges One way forward for AI in hiring involves keeping people at the center. Instead of full automation, machines suggest applicants based on match levels, offer insight summaries, track patterns in outreach, yet leave judgment calls to experienced staff. Decisions still rest with those who understand context. Tools guide, but do not replace, human oversight Working together like this mixes smart computing with human insight, so mistakes from machines drop without losing responsibility inside the company.

F. Regulatory and Legal Challenges:

Now coming into shape are rules shaped by governments and global bodies, setting how AI should be used fairly in hiring. When machines help pick job candidates, clearness matters - so does treating people equally. Someone must answer if things go wrong. Bias has no place here. Keeping personal details safe stands as a core rule. These guardrails grow stronger each year So organizations using AI for hiring need to follow job rules, avoid bias, while new AI laws take shape. As government checks grow stricter, reviewing fairness now matters more than before. Testing how systems explain choices helps meet future demands quietly rising today When groups follow clear ethics rules, trust grows in how they hire people. Good standards protect job seekers, without making the process feel cold or robotic. Fair steps mean candidates stay respected throughout hiring tech use. Public faith stays strong when companies act responsibly with smart tools.

VIII. RESEARCH GAPS

One way things have changed lately? Job applications now get sorted faster because computers help pick who might fit best. Instead of people reading every single page, software scans for keywords, ranks applicants, then suggests matches - all without stopping. What makes this possible hides inside programs trained to understand how humans write and talk. These tools borrow ideas from how brains work, spotting patterns across piles of old resumes to guess what works next time. Yet even with smart tech running tasks once done by hand, problems remain unsolved in real workplaces. Some systems favor certain backgrounds unfairly, others make choices nobody fully understands. Without fixes, trust stays low when decisions affect lives so directly. Figuring out exactly where current methods fall short helps point researchers where to look next - toward building setups that feel less like black boxes, more like open paths anyone could follow

- **Gap 1:** Limited Integration Of Fairness And Predictive Accuracy What stands out most is how few studies manage to balance accurate hiring predictions with fair outcomes. Instead of equal emphasis, current research tends to spotlight results like precision, recall, F1-score, yet overlooks whether tools treat candidates fairly across different groups. Even when methods try to include fairness, they either lower prediction quality or demand major changes to the system. Because of this trade-off, building smart hiring tools remains tricky - getting both strong accuracy and unbiased choices isn’t easily done. Moving forward, progress may depend less on chasing higher scores, more on creating approaches that lift performance and equity together for varied applicant pools.

- **Gap 2:** Lack of Standardized Benchmark Data Sets One big problem? There are no public benchmark datasets built just for studying AI in hiring. Most studies test their models on private company data or internal resume collections others can’t reach. Without shared standards, judging which algorithm works better becomes a guessing game. Results from one lab often can’t be recreated elsewhere due to locked-down sources. Hard to verify claims when everyone uses different hidden materials behind closed doors A shift toward bigger hiring data sets - anonymous, gathered responsibly - could lift studies by letting algorithms be tested side by side fairly. When these reference collections

IARJSET
ISSN (O) 2393-8021, ISSN (P) 2394-1588
International Advanced Research Journal in Science, Engineering and Technology
Impact Factor 8.066 ⵣ Peer-reviewed & Refereed journal ⵣ Vol. 13, Issue 7, July 2026
DOI: 10.17148/IARJSET.2026.13743

include people from many schools, cultures, work paths, flaws like skewed samples fade, performance widens. What matters? Fairer tests start with who's inside the data.

- **Gap 3:** Limited Understanding of How Deep Learning Models Make Decisions Even so, models like BERT, Sentence-BERT, and large language setups have gotten much better at linking jobs and skills through meaning. Yet, their inner workings stay tangled and tough to follow. Because of that, hiring staff see numbers pop up with little clue why a match was made. That lack of clarity? It chips away confidence in the tools meant to help hire people. On top of it, meeting legal rules gets harder when decisions can't be clearly explained Still, methods like SHAP and LIME show potential yet see little use in smart hiring systems. These tools add computing load; worse, they can confuse when applied to big transformer models. So new paths must explore simpler ways to explain choices - clear insights without heavy processing demands

- **Gap 4:** Insufficient Bias Checks After Launch Change shapes how hiring systems work over months and years. Workforce details shift constantly, along with what jobs demand and who applies. Most research focuses first on building models then testing them in controlled settings. Yet only a small number examine life after launch inside actual companies. Markets move, degrees transform, needs update - systems must adapt or fail Over time, if no one watches closely, machine learning tools can start making skewed choices because job markets shift and old assumptions fade. So it follows that hiring tech must keep checking its own fairness, spot imbalances as they appear, then adjust itself on the fly - staying both fair and precise from first use to last .

- **Gap 5:** Large Language Models Not Fully Used Surprisingly, big language models can now grasp meaning, follow complex ideas, even chat naturally. Still, few hiring researchers have tested these systems when picking applicants. Some explore custom job suggestions - others check how they aid interviews. A small group examines if such tools help recruiters choose better. Yet overall, real-world use remains rare despite clear possibilities One thing worth noting is how little we know about blending RAG with custom language models and smart AI helpers in hiring tools. Instead of guessing, these pieces could work together to grasp job context better. They might also cut down on made-up answers that confuse decisions. Because of this gap, using generative AI in hiring still needs deeper study. Only then can real progress take shape.

- **Gap 6:** Limited Human–AI Collaboration Right now, many AI tools in hiring care more about doing tasks fast than working alongside people. When these systems make choices alone, they can miss subtle details human recruiters notice easily - things like how someone speaks, leads, thinks differently, or fits into a team culture. Smart hiring tools of tomorrow might work better if people stay involved. Instead of letting machines decide alone, software could offer suggestions while humans make choices. This mix lets speed meet experience. One part fast processing, another part careful thinking - balance shapes fairer results. Heavy machine control often causes worry; shared control tends to ease it. Thoughtful teamwork between staff and technology may raise standards without sacrificing trust .

- **Gap 7**: Ethical and Regulatory Challenges Even as more experts agree on the need for ethical AI, only a small number of investigations lay out usable methods to meet new rules around automation and hiring laws. Consent from job applicants often gets overlooked when systems make choices without clear oversight. Clarity about how decisions are made is rare in current research on digital hiring tools. Responsibility for outcomes tends to blur when algorithms take part in screening people. Checking whether automated processes work fairly happens far too infrequently in practice. Safeguarding personal information during tech-driven recruitment still lacks strong guidance in most published work When officials keep rolling out rules for artificial intelligence, hiring tools will need to blend law-abiding features with smarter tech. Working together across fields - like coders teaming up with HR staff, lawyers, philosophers studying right and wrong, plus lawmakers - will shape how fairly these systems operate.

IX. PROPOSED GENERALIZED AI RECRUITMENT MODEL

Smart machines now help hiring teams sort through resumes and rank applicants more effectively, one way being by reading job details closely. Still, earlier parts showed current tools struggle with real-world context, carry hidden biases, often can't explain decisions clearly, rarely check if results are fair, plus miss ongoing checks once live. A fresh look comes here – combining smart text analysis, deep meaning comparisons using modern models, clear reasoning traces, balanced scoring methods into one flow people can follow step by step. This idea does not aim to wipe out what exists instead it pulls together strong ideas seen lately, fixing weak spots others leave open.

IARJSET
ISSN (O) 2393-8021, ISSN (P) 2394-1588
International Advanced Research Journal in Science, Engineering and Technology
Impact Factor 8.066 ⪽ Peer-reviewed & Refereed journal ⪽ Vol. 13, Issue 7, July 2026
DOI: 10.17148/IARJSET.2026.13743

- **Proposed System:**

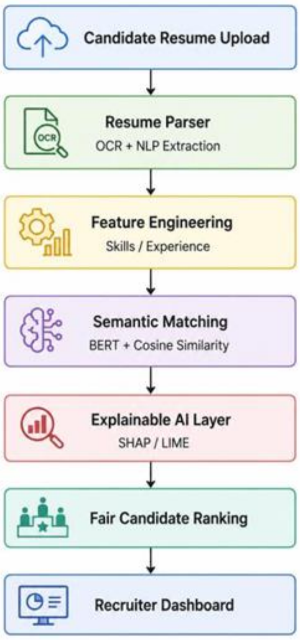

Fig .8.1: Proposed Generalized AI Recruitment Model

X. FUTURE RESEARCH DIRECTIONS

Smart computers now help pick job candidates using clever ways to sort resumes, judge applicants automatically, sometimes even match people to roles by meaning instead of keywords alone. Even though tools like language processing, machine learning, deep networks, and smart text models have pushed things forward, plenty of gaps still exist in what we know. What comes next should care less about just being right more often, more about staying fair, clear why choices are made, who answers for mistakes, how data stays safe, and whether it's used wisely. Because of this shift, future hiring tech will likely blend stronger number crunching smarts with better moral reasoning built into its core. Looking at where things are headed, mixing big language tools like GPT, Llama, or Gemini into hiring tech stands out. Instead of just matching keywords, these systems think through details - reading resumes closely, shaping candidate views, helping during interviews, suggesting roles, even backing up recruiters. Because they grasp context better, they handle messy job posts, boil down applicant backgrounds, spot skill gaps, then tailor advice more accurately. Yet before rolling them out everywhere, work must dig into ways to cut false outputs, keep facts straight, make choices steady. A new direction is taking shape around retrieval-augmented methods in hiring tools. Instead of depending only on past training data, these models pull in live company details when forming responses. Think updated rules, active job posts, field benchmarks - all fed into the output. What emerges is advice shaped by today’s context, not yesterday’s snapshots. Accuracy climbs because outdated assumptions get replaced with what actually matters now. Systems stay aligned simply by tapping fresh internal sources each time they respond One reason progress keeps moving forward is how researchers now stress fairer outcomes in machine learning. Old hiring tools typically chase high prediction scores but skip deep checks on equal treatment across groups. Instead of just boosting precision, new methods could build fairness rules right into their core logic. Bias might still slip through later if no system tracks it over time, which means safeguards need to run nonstop once software goes live. Testing at launch alone misses problems that appear months afterward One key area gaining attention is making artificial intelligence easier to understand within hiring systems built on advanced transformer models. Even though methods like SHAP and LIME offer useful insights, they demand heavy computing power - slowing down broad use in big hiring setups. Because of this, new efforts might shift toward simpler explanation tools that deliver clear reasoning quickly during live decisions, without sacrificing the strong accuracy found in today’s deep learning designs Nowadays, keeping job seeker details private is getting more attention in research. Since CVs hold personal data, upcoming hiring tools might use methods like Federated Learning, along with Differential Privacy and Secure Multi-Party Computation - so info stays protected even when models learn from it. With these techniques, companies could work together to build better selection systems, yet still keep applications locked down and aligned with global privacy rules Later on, people working alongside artificial intelligence might change how companies hire. Instead of taking over jobs, smart tools could help by sorting resumes, listing candidates, checking fair treatment, then offering suggestions. Recruiters would still make sense of situations, apply judgment, choose who gets hired. Together, human insight meets

IARJSET ISSN (O) 2393-8021, ISSN (P) 2394-1588

International Advanced Research Journal in Science, Engineering and Technology

Impact Factor 8.066 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 13, Issue 7, July 2026

DOI: 10.17148/IARJSET.2026.13743

machine speed - mistakes tied to full automation drop, responsibility stays clear within the company One step ahead, tomorrow's smart hiring tools must learn nonstop, adjusting whenever job markets shift or fresh tech appears. Because circumstances change - companies evolve, software updates - it helps when algorithms absorb what's new. Picture machines noticing rising trends: different coding styles, recent credentials, unfamiliar methods at work today. These thinking systems grow sharper over time instead of falling behind. As they adapt quietly, matching people to roles becomes less guesswork. Staying current isn't automatic; it takes design that welcomes flux. Accuracy improves simply because yesterday's data doesn't run everything anymore One big challenge sits ahead for AI hiring tools - getting them smart enough to guess right without bias. Yet they must also show how decisions happen, opening every step to view. Trust builds when people see clarity, not just results. Protection of personal data matters just as much as honest outcomes. Machines making job choices cannot skip ethics like a forgotten rule. Balance shapes reliability, linking company goals with public trust. Without these pieces fitting, progress stalls no matter the promises.

XI. CONCLUSION

Nowhere else has tech shifted hiring like artificial intelligence, reshaping who gets noticed, how they are assessed, finally brought into roles. Instead of relying only on human review, systems now parse resumes automatically thanks to tools built on machine learning, deep neural networks, plus natural language understanding. Job matches happen through meaning, not just keywords, allowing smarter suggestions that save time compared to old methods. Recruiters handle fewer routine tasks because software manages initial sorting, speeds up decisions, keeps evaluations more uniform across applicants. Productivity inside companies climbs as a result, even if challenges remain behind the scenes. One step at a time, hiring tools moved from basic tracking software to smarter systems using deep learning. Instead of just scanning resumes for keywords, today's versions grasp how skills match job needs through context. Behind this shift, models like BERT and its variants decode meaning more accurately than before. On top of that, large language networks add even greater depth to understanding fit. What stands out is not only performance but also clarity - methods such as SHAP and LIME show why decisions happen. So now, choices in hiring can be both sharper and easier to follow One thing stands clear from the studies looked at here - smart hiring tools still face tough roadblocks. Even with new methods popping up, getting accuracy without sacrificing fairness stays tricky. Biased algorithms show up again and again, along with unfair treatment across groups, worries over data misuse, murky decision-making, heavy computing loads, plus a lack of common testing rules. Transformer-powered language models are spreading fast, which only sharpens the demand for systems people can follow, question, and rely on under legal scrutiny. Trust slips when decisions feel like black box guesses. Ethics lags behind innovation, leaving gaps hard to ignore Looking at past studies closely, this work introduces a broad AI setup for hiring that brings together resume scanning, language processing, meaning comparison, clear reasoning from machines, alongside fair candidate sorting - all working inside one hiring flow. By weaving deep text comprehension with open choices and people watching the process, it tackles weak spots found in today's methods. Instead of cutting out hiring staff, the model leans toward teamwork - machines help sort and suggest, yet humans still decide who gets hired One path forward involves weaving together large language models with retrieval-augmented generation - this opens room for smarter hiring tools. Instead of chasing only speed, blending in federated learning allows data privacy to keep pace. Fairness-focused algorithms enter at this point, nudging systems toward balanced outcomes. Adaptive methods shift the rhythm, letting platforms learn without rigid reprogramming. Responsibility folds in through structured oversight, not just after problems arise but built into design. Accuracy gains mean little if decisions lack clear reasoning behind them. Trust grows when results are explainable, not just fast. So progress leans on balancing prediction strength with honesty about limits. Innovation serves everyone better when it does not leave fairness behind All things considered, AI might just speed up hiring, making it sharper and less messy. Still, these tools won't work well unless they're built with openness, fairness, clear reasoning, besides real attention to people's needs. Progress will hinge on teamwork - between tech builders, HR experts, rule makers, plus those who study right and wrong - not anytime soon, but steadily, carefully

OVERALL SUMMARY

Looking closely at today's hiring tools, this paper explores how artificial intelligence is reshaping the way resumes get reviewed and matched to jobs. Old-school hiring tends to drag on, demands heavy manual effort, while quietly letting personal biases slip through. Instead of relying only on people, companies now lean on smart tech - like NLP, ML, DL, and advanced setups such as BERT and Sentence-BERT - to dig into résumés with sharper accuracy. These systems go beyond keywords, grasping meaning, aligning skills more thoughtfully with roles, then sorting applicants in smarter ways. On top of that, explainable AI steps in - not just deciding who fits, but showing why certain choices emerge, which helps make automated decisions clearer and fairer across the board. Looking at new studies side by side reveals what works - and where current AI hiring tools fall short. Gaps stand out when it comes to fairness, data privacy, clear reasoning behind choices, and missing uniform ways to measure success. One idea gaining ground mixes resume scanning with

International Advanced Research Journal in Science, Engineering and Technology
Impact Factor 8.066 ⪤ Peer-reviewed & Refereed journal ⪤ Vol. 13, Issue 7, July 2026
DOI: 10.17148/IARJSET.2026.13743

meaning-based job fit scoring, balanced rankings, and open explanations of decisions. Efficiency and smarter picks seem within reach if machines help sort applicants. Still, trust grows only when ethics lead the way, hidden biases get checked often, people stay involved, and how things work stays visible. Systems built like that might actually serve today's workplaces well.

REFERENCES

[1]. E. Albaroudi, T. Mansouri, and A. Alameer, "A Comprehensive Review of AI Techniques for Addressing Algorithmic Bias in Job Hiring," AI Journal, vol. 5, pp. 383 404, 2024.
[2]. N. Akhtar, S. Rabbani, H. Rabbani, and S. Kumar, "AI-Driven Intelligent Resume Recommendation Engine," International Journal of Scientific Research in Science, Engineering and Technology, vol. 12, no. 3, pp. 1141–1155, 2025.
[3]. R. Muthu Kumar, "AI-Based Resume Analyzer Using NLP and Machine Learning," International Journal of Innovative Research in Technology, vol. 12, no. 1, pp. 1215–1220, 2025.
[4]. M.-H. Ajjam and H. S. Al-Raweshidy, "AI-driven Semantic Similarity-based Job Matching Framework for Recruitment Systems," Information Sciences, vol. 724, 2026.
[5]. P. Desai, N. Khankure, and S. Kolekar, "AI-Powered Resume Screening and Job Matching System," International Journal of Sciences and Innovation Engineering, vol. 2, no. 11, pp. 132–136, 2025.
[6]. K. S. Yadav et al., "AI-Powered Resume Screening and Job Matching System for Intelligent Career Guidance and Recruitment Optimization," International Journal on Science and Technology, vol. 16, no. 4, 2025
[7]. S. Yadav, S. Ursal, and S. Tate, "Resume Analysis Using NLP and ATS Algorithm," International Journal of Latest Technology in Engineering, Management & Applied Science, vol. 14, no. 4, pp. 761–768, 2025.
[8]. S. B. Anusha et al., "Resume Screening & Job Matching System: A Machine Learning Approach," Journal of Artificial Intelligence and Computational Technology, vol. 2, no. 2, 2025.
[9]. N. Alsubaie and N. Aleisa, "Mitigating Bias in AI Model using Explainable AI in Terms of Hiring Process in the Industry," IEEE Access, 2025.
[10]. S. Dogra et al., "Smart Resume Screening and Matching System," National Conference on AI, IoT & Data-Driven Transformation, 2025.
[11]. S. M. A. Idid et al., "AI-Powered Resume Crafting and Screening," International Journal on Perceptive and Cognitive Computing, vol. 12, no. 1, 2026.
[12]. P. V. J., "Resume Parsing and Job Prediction Using CNN and NLP," Journal of Intelligent Systems, 2024.
[13]. D. F. Mujtaba and N. R. Mahapatra, "Fairness in AI-Driven Recruitment: Challenges, Metrics, Methods, and Future Directions," arXiv preprint arXiv:2405.19699, 2025.
[14]. K. Wilson and A. Caliskan, "Gender, Race, and Intersectional Bias in Resume Screening via Language Model Retrieval," AAAI Conference on Artificial Intelligence, 2024.
[15]. R. K. Magham, "Mitigating Bias in AI-Driven Recruitment: The Role of Explainable Machine Learning (XAI)," International Journal of Scientific Research in Computer Science, Engineering and Information Technology, vol. 10, no. 5, pp. 461–469, 2024.
[16]. G. Sravan Kumar, K. Varshitha, K. Keerthi, and N. Vikesh, "Resume Screening Automation with NLP Techniques," International Journal of Innovative Science and Research Technology, vol. 10, no. 5, pp. 3294–3298, 2025.
[17]. K. Khelkhal and D. Lanasri, "SMART-HIRING: An Explainable End-to-End Pipeline for CV Information Extraction and Job Matching," arXiv preprint arXiv:2511.02537, 2025.
[18]. J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," Proceedings of NAACL-HLT, 2019.
[19]. N. Reimers and I. Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT Networks," Proceedings of EMNLP, 2019.
[20]. S. Lundberg and S. Lee, "A Unified Approach to Interpreting Model Predictions," Proceedings of NeurIPS, 2017.
[21]. P. Resnick and H. Varian, "Recommender Systems," Communications of the ACM, vol. 40, no. 3, pp. 56–58, 1997.
[22]. C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel, "Fairness Through Awareness," Proceedings of ITCS, 2012.
[23]. R. Zemel, Y. Wu, K. Swersky, T. Pitassi, and C. Dwork, "Learning Fair Representations," Proceedings of ICML, 2013.
[24]. European Commission, "Ethics Guidelines for Trustworthy Artificial Intelligence," High-Level Expert Group on AI, 2019.
[25]. OECD, "Principles on Artificial Intelligence," OECD AI Policy Observatory, 2024.