

Intelligent Job Matching and Allocation Using Machine Learning, Language Models, and Strategy proof Mechanisms: A Review

¹Hande Pooja Vinod, ² Prof.Bangar Abhishek P., ³Prof.Karad Vandana A., ⁴Prof.Bhosale Sachin B.

^{1,2,3,4}Department of Computer Engineering / Jaihind College of Engineering / SPPU, Pune, India

Abstract — Job matching is increasingly supported by computational systems that estimate employment opportunities, interpret unstructured career information, and allocate scarce positions or resources. This review synthesises three supplied research papers that address complementary parts of this problem: large-scale machine learning for labour-market matching, strategyproof matching when project capacities depend on resource allocation, and CareerBERT for semantic resume-to-job recommendation. The first study analyses approximately 55 million German employment spells from 2012–2018 and reports that random forests outperform logit, ordinary least squares, and k-nearest neighbours for occupation-level prediction. The second formalises a student-project-resource problem and shows that strategyproofness cannot be combined with fairness and weak efficiency in full generality; it introduces Sample and Deferred Acceptance (SDA) mechanisms to balance these objectives. The third combines ESCO occupational taxonomy entries with EURES job advertisements and fine-tunes a BERT-based Siamese network using Multiple Negatives Ranking loss to place resumes and jobs in a shared semantic embedding space. Across the papers, matching quality depends on the representation of preferences and profiles, the scale and quality of data, the way capacities are modelled, and the objectives used to evaluate outcomes. Predictive accuracy and semantic relevance can improve recommendations, but they do not by themselves guarantee fairness, truthful reporting, or efficient resource use. The review therefore proposes a layered design that combines data-driven ranking with transparent allocation rules, human oversight, and explicit auditing of outcomes. Since the review is limited to three supplied studies with different tasks and evaluation settings, reported metrics are compared descriptively rather than treated as directly equivalent.

Keywords— Job Matching; Labour Market; Machine Learning; Random Forest; CareerBERT; BERT; Natural Language Processing; Strategyproofness; Resource Allocation; Deferred Acceptance; Fairness.

I. INTRODUCTION

Matching people to suitable work, training, projects, or educational opportunities is a central coordination problem in modern labour markets and public services. The task appears simple when each candidate has a short list of preferences and every position has a known capacity. Real systems are more complicated. Candidate profiles are heterogeneous, job descriptions are written in varied language, occupational categories may number in the hundreds, and the capacity of a project or service may depend on resources that must themselves be allocated. A useful matching system must therefore address both prediction—estimating which opportunities are suitable—and allocation—deciding which assignments can be made while respecting preferences and constraints.

The three studies considered here approach these challenges from different angles. Mühlbauer and Weber [1] frame matching as the prediction of the occupational field entered in an employment spell. Their large administrative dataset allows

algorithms to learn patterns across employment histories and to produce ranked occupation recommendations. Rosenberger et al. [2] focus on semantic matching between resumes and general job descriptions, using domain-adapted language modelling to capture meaning beyond exact keyword overlap. Liu, Yahiro, and Yokoo [3] study a formal allocation problem in which students prefer projects, projects rank students, and indivisible resources determine project capacities. Their analysis highlights that a desirable assignment cannot be assessed by predictive accuracy alone: strategic incentives, fairness, and efficiency matter as well.

These contributions are related but not interchangeable. A statistical model can estimate the probability that a person enters an occupation without proving that a recommendation is fair. A semantic model can identify a relevant job title from a resume without determining how a limited number of places should be distributed. A strategyproof mechanism can protect incentives under its assumptions but may not itself understand free-form resumes or predict future transitions. Integrating these perspectives is important for decision-support systems

intended for real employment services, recruitment, education, or training allocation.

This review has four objectives: (i) summarise the problem settings, methods, and reported findings of the three papers; (ii) compare their data representations and evaluation approaches; (iii) identify complementary strengths and limitations; and (iv) propose an integrated framework and research agenda for responsible intelligent matching. The review does not claim to be a systematic review of the entire field. It is a structured narrative synthesis of the three source papers supplied for this task.

II. REVIEW METHODOLOGY

The review corpus consists of three full research articles provided as source PDFs. The papers were selected by the user and are treated as the primary evidence base: (1) “Machine learning for labor market matching” by Mühlbauer and Weber; (2) “Strategyproof Mechanism for Two-Sided Matching with Resource Allocation” by Liu, Yahiro, and Yokoo; and (3) “CareerBERT: Matching resumes to ESCO jobs in a shared embedding space for generic job recommendations” by Rosenberger and colleagues.

Each paper was read for its research objective, model or mechanism, data, evaluation design, reported results, limitations, and practical implications.

The synthesis uses a comparative narrative approach. First, each paper is summarised independently to preserve its own terminology and claims. Second, a common comparison framework is applied across five dimensions: input representation, matching objective, technical method, evaluation evidence, and safeguards or constraints.

Third, the papers are considered jointly to identify a possible system architecture. Because the studies use different units, tasks, datasets, and metrics, their numerical results are not pooled and no cross-paper performance ranking is calculated. The review also avoids inferring causal impact where the source papers report prediction, simulation, or experimental comparisons rather than a real-world intervention.

Scope note: The review is grounded in the three supplied articles. Statements about their findings are attributed to the relevant paper; the integrated architecture and recommendations in later sections are the authors’ synthesis rather than a system evaluated by any of the three studies.

III. BACKGROUND AND PROBLEM FORMULATION

1. Predictive matching

Predictive matching estimates how likely a candidate is to enter, succeed in, or be selected for a particular occupational category. In the labour-market study [1], the target is the occupation observed in an employment spell. The model uses individual and job-related characteristics from administrative employment biographies to estimate occupation-specific probabilities. A recommendation interface can then rank occupations by predicted probability. The central technical challenge is high-dimensional, imbalanced classification: the paper considers up to approximately 700 occupational categories, whose observed frequencies differ substantially.

2. Semantic resume-to-job matching

Resume-to-job matching is a text representation problem. A resume may describe skills and experience using phrases that differ from the terminology used in a job title or formal occupation taxonomy. Exact word overlap is therefore not always a reliable indicator of relevance. CareerBERT [2] learns a shared embedding space in which semantically related resumes and jobs should be close together. The use of the European Skills, Competences, Qualifications and Occupations (ESCO) taxonomy provides standardised occupational concepts, while EURES advertisements contribute real-world descriptions and current labour-market language.

3. Preference-based allocation

The Student-Project-Resource (SPR) model [3] extends ordinary two-sided matching. Students rank projects and projects rank students, but project capacity is not fixed independently: it depends on indivisible resources allocated to each project. A resource can be assigned to only one compatible project, even if its capacity could otherwise be divided. Separating resource allocation from student matching can produce a poor result when the demand assumptions used to set capacities are inaccurate. The SPR framework therefore treats allocation and matching as one coupled problem.

Together, these formulations show that “matching” refers to several related but distinct tasks. Prediction estimates likely outcomes from historical data; semantic retrieval measures textual or profile relevance; and a matching mechanism allocates scarce opportunities according to preferences and constraints. A deployable service may need all three, but should not treat any one as a substitute for the others.

IV. LITERATURE REVIEW: THREE CORE STUDIES

1. Machine learning for labour-market matching

Mühlbauer and Weber [1] develop a large-scale framework for estimating individual matching probabilities across occupational groups. Their data cover approximately 55 million employment spells and represent around six percent of the German workforce over 2012–2018. The longitudinal administrative records include personal characteristics and job-related information. The objective is to use these histories to produce occupation-level recommendations that can support public employment service caseworkers and job seekers.

The authors compare ordinary least squares (OLS), logit, k-nearest neighbours (kNN), and random forest (RF). Random forests aggregate decision trees trained on bootstrap samples and introduce random feature selection at tree splits. This structure allows nonlinear relationships and interactions to be represented without requiring the analyst to specify every functional form in advance.

The study reports that RF consistently achieves the highest predictive performance and that its out-of-sample classification error is approximately 30 percent lower than that of OLS. The authors also find that the model benefits from using all employment spells rather than restricting estimation to unemployment-to-employment transitions.

The paper's policy interpretation is deliberately potential rather than a demonstrated intervention effect: its results suggest that machine-learning-based matching could hypothetically reduce unemployment by up to 0.3 percentage points. This estimate should not be read as proof that deploying the model will produce that reduction. Actual outcomes would depend on whether recommendations change search behaviour, whether suitable vacancies exist, and how caseworkers and job seekers use the ranked suggestions.

A major contribution is the fit between the algorithm and the data structure. Administrative labour-market data are large, longitudinal, heterogeneous, and often imbalanced. The paper argues that conventional models can be limited by restrictive functional forms or computational difficulty, whereas RF can continue to exploit additional observations. At the same time, the authors explicitly recognise potential discrimination when historical patterns and sensitive proxies are learned by algorithms. Their discussion emphasises examining discriminatory influence rather than assuming that predictive performance alone is sufficient.

2. Strategyproof matching with resource allocation

Liu, Yahiro, and Yokoo [3] define the SPR problem, where a set of students, projects, and indivisible resources must be coordinated. Students and projects have strict preferences over one another, while resources have capacities and compatibility lists that restrict the projects to which they may be allocated. A project's effective capacity is induced by the resources it receives. The authors motivate the problem with settings such as student-project allocation, school choice with flexible capacity, and nursery-school staffing, where the resources needed to expand capacity may not be known precisely before preferences are revealed.

The paper demonstrates an impossibility result: no strategyproof mechanism can satisfy fairness and the specified weak efficiency requirements simultaneously in the general SPR setting. This result is important because it shows that the design problem is not solved merely by selecting a better optimisation algorithm. The desired properties can conflict mathematically. The authors also discuss limitations of existing approaches, including Serial Dictatorship (SD), Artificial Caps Deferred Acceptance (ACDA), and Adaptive Deferred Acceptance (ADA).

To address the trade-off, the authors introduce Sample and Deferred Acceptance (SDA), a family of mechanisms. Students are split into sampled and regular groups. Sampled students are assigned using serial dictatorship; their preferences inform resource allocation; then regular students are matched using Deferred Acceptance (DA) under the resulting capacities. The paper studies variants including SDA with Voting (SDA-V), a variant using a more sophisticated voting procedure (SDA-V*), and SDA with Simulation (SDA-S). Experimental results indicate that SDA variants can balance fairness and efficiency, and SDA-S performs better than other SDA variants when students are divided into multiple preference types.

Unlike a predictive model, SDA is a rule for making assignments given preferences and constraints. Its value lies in incentive and allocation properties, not in predicting whether a particular occupation is likely to be entered. Its limitations also follow from its formal setting: the mechanism relies on the model's assumptions about preferences, resources, and feasible assignments, and the experiments do not establish that one mechanism will dominate in every real-world matching market.

3. CareerBERT for resume-to-job recommendations

Rosenberger et al. [2] introduce CareerBERT to support job seekers and career counsellors with generic job recommendations based on resumes. Rather than limiting

recommendations to a fixed set of advertisements, the system represents general job titles through ESCO and augments this taxonomy with EURES job advertisements. This combination is intended to provide a standardised yet current representation of occupations and their descriptions.

The system has two main components: a domain-adapted language model and a resume-job matching module. The authors compare base models, including GBERT and jobGBERT, with and without domain-adaptive pre-training. The selected model is fine-tuned for matching using a Siamese network architecture and Multiple Negatives Ranking (MNR) loss. Job titles are paired with corresponding skills, descriptions, or synonyms. The training objective draws matching pairs closer in embedding space and separates non-matching examples through a ranking-based loss. Cosine similarity can then be used to retrieve job concepts relevant to a resume.

CareerBERT is evaluated in two stages: an application-grounded evaluation using EURES job advertisements and a human-grounded evaluation involving real-world resumes and feedback from human-resources experts. The authors report that CareerBERT outperforms traditional and state-of-the-art embedding approaches in their evaluations and shows robust effectiveness in expert assessment. The work also uses

preprocessing and relevance-filtering techniques to reduce noise, and the authors make experimental code and model artefacts publicly available to support reproducibility.

The study's central strength is semantic generalisation: a candidate's experience may be relevant to an occupation even when resume wording does not exactly match a job title. However, semantic similarity does not guarantee availability of vacancies, suitability under legal or credential requirements, or fairness across demographic groups. Human evaluation is valuable, but the recommendation model should still be assessed for ranking quality, coverage, calibration, and potential bias before being used to make consequential decisions.

V. COMPARATIVE ANALYSIS

Table 1 summarises the three studies. Their outcomes should be interpreted in relation to their own objectives: the first reports predictive classification performance, the second examines mechanism properties and experimental trade-offs, and the third evaluates semantic retrieval and expert relevance. No single metric allows a direct ranking across all three.

Table 1. Comparative summary of the three reviewed studies.

Dimension	Mühlbauer & Weber [1]	Liu et al. [3]	Rosenberger et al. [2]
Primary problem	Predict occupation entered in an employment spell	Match students and projects when resources determine capacity	Recommend general jobs from resume text
Main input	Longitudinal administrative employment biographies	Preferences, resource capacities, compatibility constraints	Resumes, ESCO taxonomy, EURES advertisements
Core method	OLS, logit, kNN, random forest	Strategyproof mechanisms; SDA family with DA	Domain-adapted BERT; Siamese network; MNR loss
Reported result	RF best; about 30% lower out-ofsample error than OLS	Impossibility result; SDA balances fairness and efficiency	CareerBERT outperforms comparison embedding approaches
Evaluation	Large-scale held-out prediction	Formal properties and experiments/simulation	Advertisement-based and HR expertgrounded evaluation
Key strength	Scale and nonlinear pattern learning	Explicit incentive and allocation reasoning	Contextual semantic understanding of text

Main caution	Historical bias; prediction is not causal impact	Fairness and efficiency can conflict; assumptions matter	Semantic relevance does not ensure fairness or feasibility
--------------	--	--	--

1. Representation of people and opportunities

The studies differ most visibly in how they represent a person and an opportunity. Administrative records in [1] encode employment histories and structured characteristics. This representation is powerful for learning from large numbers of observed transitions, but it reflects what has been recorded historically and may underrepresent new occupations or atypical career paths. CareerBERT [2] uses free-form text and standardised occupational descriptions, enabling a richer treatment of skills expressed in different words. Yet text can be incomplete, inflated, or inconsistently written, and the model's outputs depend on the corpus used for training. The SPR model [3] instead assumes preferences are available as ordered lists; it does not solve the natural-language extraction problem. In a practical system, preference elicitation and profile interpretation would precede or accompany the allocation mechanism.

2. Objectives and what counts as a good match

The first study optimises predictive accuracy for occupation categories, the second reasons about fairness, strategyproofness, and efficiency, and the third emphasises semantic ranking relevance. These objectives overlap but can disagree. The most probable occupation may not be the person's preferred next step. The semantically closest job may not have an available vacancy or may require credentials that the resume does not establish. A mechanism that protects truthful preference revelation may assign a candidate to a feasible but less semantically similar option. Consequently, a system should state whether its output is a prediction, a recommendation, or a binding allocation, and evaluate each function against its own criteria.

3. Scale, interpretability, and deployment

The administrative-data model [1] demonstrates the benefit of scale and flexible nonlinear modelling. CareerBERT [2] offers semantic representation and retrieval suitable for large collections of job concepts, while SDA [3] provides an explicit procedural structure for allocating constrained capacity. These components impose different deployment requirements. Prediction and retrieval need monitoring for data drift and ranking quality; allocation needs a clear record of preferences, resource compatibility, tie-breaking rules, and feasible outcomes. In all cases, staff and users need explanations that

are meaningful enough to challenge errors and correct missing information.

VI. KEY FINDINGS AND CROSS-STUDY SYNTHESIS

Finding 1: The quality of the representation is decisive. The studies show three complementary forms of representation: structured employment histories, contextual text embeddings, and explicit preference/resource models. Each captures information that the others do not. Structured histories reflect realised transitions; language embeddings capture semantic similarity; and formal preference models represent allocation constraints and strategic behaviour.

Finding 2: Better prediction is not the same as a better allocation. Random forest's lower classification error [1] is a result about predicting observed occupations, not a proof of fairness or welfare improvement. CareerBERT's semantic ranking [2] supports counsellors but does not by itself allocate scarce positions. The impossibility result in [3] demonstrates that even a well-designed mechanism may face unavoidable trade-offs among desired properties. A complete system needs both relevance and allocation safeguards.

Finding 3: Fairness must be designed and measured explicitly. Historical employment data may encode unequal access to occupations, while text-based systems may be affected by differences in resume style, language, or career opportunity. Strategyproofness addresses incentives to misreport preferences under the formal model, but it does not automatically remove all forms of distributive or demographic bias. Fairness therefore requires separate tests, clear definitions of protected outcomes, and ongoing monitoring rather than reliance on a single algorithmic property.

Finding 4: Human oversight remains important. The CareerBERT study includes HR expert evaluation [2], while the labour-market study positions recommendations as decision support for public employment service caseworkers [1]. These settings suggest that ranked recommendations are most defensible when users can review and override them, record the reason for an override, and provide feedback. Human involvement should not be treated as a substitute for model validation, because reviewers may also inherit existing biases.

Finding 5: Evaluation should follow the intended use. Large-scale held-out classification, embedding-based retrieval tests, expert relevance judgements, formal mechanism properties, and allocation simulations answer different questions. Future systems should report not only average scores but also missed opportunities, coverage for smaller groups, calibration, stability under distribution shift, feasibility, and the effect of recommendations on real outcomes where it can be ethically and reliably measured.

VII. PROPOSED INTEGRATED FRAMEWORK

Based on the complementary findings, this review proposes a conceptual architecture with six layers. This is a synthesis for future investigation, not an architecture implemented or validated by the three source papers.

1. Layer 1: Data ingestion and governance

The system would ingest structured employment histories, resume text, standard occupational taxonomies, job descriptions, vacancies, and explicit user preferences. Provenance, timestamps, consent, data minimisation, and retention rules should be recorded. Administrative and text datasets should be quality-checked for missing values, duplicate records, stale job descriptions, and inconsistent occupation labels. Where historical data are used, the system should document which populations and time periods are represented.

2. Layer 2: Profile and occupation representation

A structured feature pipeline can represent stable employment and qualification information, while a domain-adapted language model can encode resumes and job descriptions. An occupational taxonomy such as ESCO can act as a shared vocabulary connecting free text with standard job concepts. The system should retain source evidence for extracted skills and avoid interpreting semantic similarity as proof that a candidate possesses a mandatory credential.

3. Layer 3: Candidate-specific prediction and retrieval

A predictive model, such as the random-forest approach studied in [1], can estimate likely occupational transitions from structured histories. A semantic retrieval model, such as CareerBERT [2], can find related job concepts from resume content. These outputs should be combined cautiously—for example, by presenting separate signals or using a validated ranking method—rather than assuming that scores from different models are calibrated or directly comparable. Results

should include reasons or supporting profile features where feasible.

4. Layer 4: Feasibility and allocation

Before a recommendation becomes an assignment, the system should check vacancy availability, eligibility, location, working conditions, and capacity. Where allocation involves multiple stakeholders and resources, a formal mechanism can encode preferences and constraints. The SPR findings [3] indicate that mechanism designers should specify which properties are prioritised and disclose unavoidable trade-offs. Strategyproofness, fairness, nonwastefulness, and efficiency should be treated as separate properties with explicit tests.

5. Layer 5: Human review and user control

Career counsellors and users should be able to inspect recommendations, correct profile information, state preferences, and reject unsuitable options. The interface should distinguish “predicted likely occupation” from “recommended job” and from “allocated position.” Override reasons and user feedback can support system improvement, but should be reviewed to ensure they do not simply reinforce existing patterns of exclusion.

6. Layer 6: Auditing and continuous evaluation

Evaluation should cover retrieval relevance, top-k recommendation quality, calibration, vacancy coverage, allocation feasibility, user satisfaction, and longer-term employment outcomes. Where legally and ethically appropriate, results should be audited across relevant population groups and career stages. The model should be monitored for drift as occupations, skills, and vacancy distributions change.

An independent audit trail should record data versions, model versions, ranking logic, mechanism settings, and human overrides.

1. Governed data Employment records • resumes • vacancies • preferences
2. Shared representation Structured features + occupation taxonomy + text embeddings
3. Predict and retrieve Occupation probabilities + semantic job ranking
4. Feasibility and allocation Eligibility • capacities • preferences • mechanism rules
5. Human decision support

Counsellor review • user correction • transparent explanations
6. Audit and feedback Fairness • quality • outcomes • drift monitoring

Fig. 1. Proposed layered workflow for an intelligent matching and allocation service.

Challenges and Research Gaps

Dataset shift and representativeness

Employment patterns and job language change over time. Administrative histories may reflect past hiring practices, while a resume corpus may overrepresent certain occupations, languages, or levels of digital literacy. Taxonomies also need maintenance as job titles and skills evolve. A future system should measure time-based drift, test on newer cohorts, and document where its recommendations have limited evidence.

Fairness, privacy, and explainability

The reviewed studies raise distinct governance questions. The administrative-data approach uses rich personal histories, creating a need for privacy protection and scrutiny of proxy variables [1]. A text model may infer signals from writing style or career gaps that are not legitimate measures of job suitability [2]. A mechanism may be strategyproof under formal assumptions but still distribute opportunities unequally under the chosen policy objective [3]. Research should therefore combine privacy-preserving data practice, group-wise error analysis, contestable explanations, and mechanism-level audits.

Evaluation beyond offline metrics

Offline classification and retrieval results are useful but incomplete. A recommendation can be technically relevant and still fail because no vacancy exists, the candidate does not want the job, or the recommendation discourages exploration of a new career path. Controlled deployment studies should assess whether decision support improves search breadth, interview opportunities, placement duration, retention, and user trust. Where randomised experiments are not feasible, carefully designed quasi-experimental approaches may help distinguish model effects from changes in labour-market conditions.

Integration of ranking and mechanisms

A major open challenge is how to connect a probabilistic or semantic ranking score to an allocation mechanism. One option is to use model scores only to propose a candidate set, while preserving explicit preferences and capacity constraints for the final assignment. Another is to present multiple recommendations with different trade-offs rather than

collapsing them into one opaque score. Future research should examine how such interfaces affect truthful preference reporting and whether ranking assistance changes the strategic environment assumed by formal mechanisms.

Robustness and accountability

The three papers do not jointly establish a complete production-ready system. Important questions remain about score calibration, adversarial or misleading resumes, missing preference information, tie-breaking, appeals, and accountability when the model and mechanism recommend different outcomes. Operational research should specify who is responsible for the final decision, how errors can be corrected, and how users can challenge an assignment.

Recommendations for Practice

For public employment services and career consultants, machine-learning predictions should be used as an additional source of evidence, not as a substitute for an interview or the job seeker's own goals. Ranked lists should include diverse but plausible alternatives, and the interface should make clear why an occupation appears. The reported performance advantage of random forests [1] supports their evaluation for large administrative datasets, but deployment should follow local validation, bias testing, and impact assessment.

For resume-based recommendation products, domain adaptation and shared embeddings offer a way to connect varied resume language to standard occupational concepts [2]. Practitioners should maintain the taxonomy and advertisement corpus, validate recommendations with domain experts, and separate semantic relevance from eligibility and vacancy availability. A strong embedding score should not be presented as a hiring probability unless the model was trained and validated for that specific target.

For systems allocating scarce places, the allocation rule should be selected before deployment and its objectives should be documented. The SPR impossibility result [3] makes it important to explain which fairness and efficiency properties are guaranteed, which are approximate, and which cannot be jointly achieved under the model. Mechanism simulation should use realistic preference distributions and resource constraints, and the final system should retain auditable records of capacities and assignments.

Across all settings, meaningful human oversight requires more than a nominal approval step. Reviewers need enough information to understand recommendations, time to challenge errors, and a process for recording decisions. Users should be

able to correct inaccurate profile data and state constraints that are not visible in historical records. Institutions should regularly assess whether the system expands opportunity or merely reproduces existing patterns.

allocation,” *Artificial Intelligence*, vol. 316, art. no. 103855, 2023. doi: 10.1016/j.artint.2023.103855.

VIII. CONCLUSION

This review has examined three complementary approaches to intelligent matching. Mühlbauer and Weber [1] show that random forests can exploit large administrative employment histories to improve occupation-level prediction over traditional statistical baselines. Liu, Yahiro, and Yokoo [3] demonstrate that matching with endogenous capacities creates fundamental trade-offs between strategyproofness, fairness, and efficiency, and propose the SDA family as a practical compromise. Rosenberger et al. [2] show that a domain-adapted BERT model can match resumes to general occupational concepts in a shared embedding space and outperform comparison embedding methods in their evaluations.

The main conclusion is that reliable matching requires more than a single high-performing algorithm. Historical prediction, semantic understanding, and allocation rules solve different parts of the problem. A responsible system should combine these capabilities only through a clearly specified workflow, keep recommendations distinct from binding decisions, and evaluate predictive quality, relevance, feasibility, fairness, and user outcomes separately. The proposed layered framework offers a basis for future work, but it requires empirical validation in real service settings. Future research should focus on cross-dataset robustness, transparent integration of ranking with strategyproof mechanisms, rigorous fairness evaluation, and longitudinal evidence of whether intelligent recommendations improve employment and allocation outcomes.

REFERENCES

1. S. Mühlbauer and E. Weber, “Machine learning for labor market matching,” *Machine Learning with Applications*, vol. 23, art. no. 100861, 2026. doi: 10.1016/j.mlwa.2026.100861.
2. J. Rosenberger, L. Wolfrum, S. Weinzierl, M. Kraus, and P. Zschech, “CareerBERT: Matching resumes to ESCO jobs in a shared embeddingspace for generic job recommendations,” *Expert Systems with Applications*, vol. 275, art. no. 127043, 2025. doi: 10.1016/j.eswa.2025.127043.
3. K.-g. Liu, K. Yahiro, and M. Yokoo, “Strategyproof mechanism for two-sided matching with resource