Domain-First Intelligence: An AI-Driven Pre-Screening Framework for University Placement Cells to Mitigate CGPA-Centric Filtering Bias in Campus Recruitment

Dev Sanghavi[1], Khushali Amreliya[2], Shubham Sarkar[3], Dr. Sunil Mahajan[4]
[1,2,3]*Student, Master of Science (M.Sc.) Information Technology, MIT Arts, Commerce and Science College, Alandi (D), Pune, Maharashtra, India*
[4]*Associate Professor, Department of Electronics, MIT Arts, Commerce and Science College, Alandi (D), Pune, Maharashtra, India*
doi.org/10.64643/ijirt.208432-459

Abstract*—Placement in most institutions of engineering and technology still depends on an absolute cumulative grade point average (CGPA) cut-off, typically CGPA ≥ 7.0, to shortlist which students even qualify for presentation to potential employers. In effect, a significant number of "average-CGPA" students are precluded from evaluation of their job-related competencies even before such evaluation takes place, despite CGPA being a poor and error-prone measure of job competence. Following recent findings of discrimination in traditional and artificial intelligence-assisted hiring procedures, we suggest here a competency-focused, fairness audited placement process wherein students are rated through competency assessment relevant to the specific job, effort-based bonus rating (internship, projects, certification), and an optional process of blind screening prior to candidate review by a recruiting company. Due to the fact that the pipeline hasn't been developed yet, simulation is used to validate its design, where the population consists of 1,000 students and is calibrated to a realistic CGPA range (3.5–9.9) divided according to the criteria of gender, prior exposure to the field and first-generation status; and where the competency-based pipeline is run on this population along with the CGPA-only pipeline. The simulation results show that the competency-based pipeline provides an opportunity for assessment to 974 out of 1,000 students (97.4%), compared to 311 (31.1%) in the case of the CGPA-only pipeline, and increases the portion of average-CGPA students in the assessed population from 0% to 97.5%; in addition, the pass rate stays over 93% for all subgroups of the population examined, and pass rate grows monotonically depending on the degree of prior exposure to the field (93.6%→100%), which means that the test continues to discriminate between students based on their skills but not on their backgrounds.

***Index Terms*—campus placement; algorithmic fairness; blind screening; competency-based assessment; AI-assisted recruitment; bias mitigation**

I. INTRODUCTION

Campus placements can be considered the most important phase in the life of engineering/technology students, because it is through this process that the entire effort put into study for years finally gets transformed into a job opportunity. TPOs find themselves at the very heart of campus placements, balancing between hundreds of students and scores of recruiting firms, each having enough time to interview just a small number of the total candidates available to them. In order to deal with this problem, most educational institutions have come up with the idea of an eligibility test.
CGPA acts as that filter in almost all Indian colleges. There is a widely accepted criteria of CGPA ≥ 7.0 which is considered while making a decision on who should be allowed to take the test/interview of the company. Students falling below the criteria - despite being good at communication skills, programming, or having an experience of working on a project - are disqualified before testing any job-specific skill of theirs. It is a strict criterion since a student with CGPA 6.9 will be treated in the same way as the student with CGPA 3.5.

It is indeed possible to argue that there is sufficient reason why this issue should be viewed in terms of fairness rather than strictly an academic one. There is mounting evidence about the problem of bias in both artificial intelligence-driven and traditional recruitment practices which suggests that the hiring process – whether manual, rules-based or large language model (LLM)-based – is biased against certain candidates irrespective of their suitability for a position [1], [6], [7]. In research concerning recruitment beyond algorithms, there is also ample evidence that the selection process is influenced by political and structural factors, including clientelistic relations, ethnicity/regional bias and institutional incentives [4] which impact the hiring process in a manner which cannot be attributed to any "biased algorithm".

All previous approaches to campus placement modernization revolve around logistics — automating application management, scheduling of interviews, generation of shortlists — while completely neglecting issues of equity. As described in the literature on intelligent placement portals, such solutions optimize the TPO-student interaction process [10], [11], but do not change the main criterion of selection which remains CGPA or some other equally crude and potentially biased one and does not implement any auditable mechanism for ensuring fairness in the process. Meanwhile, the literature on corporate recruitment offers methods of reducing the impact of biases and in particular blind screening — the technique that strips all identity-related information from the resume or CV before any human/algo-based review takes place [2], [3]. Such technique has not been applied, to our best knowledge, to the problem of campus placement due to different nature of the recruiter community, applicant pool, and constraints in the environment. It means that there is a gap to fill — there is no system that implements a combination of the following: (a) competency-based replacement for CGPA; (b) an optional stage of blind screening proved in corporate environments; (c) fairness auditing.

There are three main contributions in this paper. First, we introduce a fair placement framework consisting of four stages, namely the competency test, effort-based scoring, blind screening, and an overall placement scheme without a hard CGPA threshold. Second, since the framework is yet to be put into practice, its effectiveness is validated via a controlled simulation using a synthetic sample size of 1,000 students with comparison between the current CGPA filter and the new framework. Third, we present a multidimensional fairness audit based on the following attributes: gender, CGPA groups, regions, economic status proxy, first generation, and previous domain exposure. Furthermore, to ensure the validity of the test proposed by the framework, we conduct an exposure-based pass rate gradient check.

II. LITERATURE REVIEW

The present literature review is structured based on the themes that form the basis of the framework in question: (A) Bias in AI-assisted recruiting, (B) Blind screening as a mitigating mechanism, (C) Factors influencing biased selection decisions, (D) Challenges to the idea of algorithmic neutrality, (E) Measures of fairness and evaluation tools, (F) Quality of AI-assisted test items, and (G) Existing campus recruitment systems.

A. Bias in AI-assisted recruiting

Systematic review of bias and fairness of LLM-based recruiting identifies a body of research indicating that language-model-based recruiting systems are susceptible to encoding biases that exist in their training datasets and resume corpora that are used for their fine-tuning or prompting [1]. Related empirical research specifically analyzing hiring biases in large language models indicates that model output is systematically variable depending on candidate attributes which are unrelated to his or her job performance [6]. Another relevant study testing LLMs in the area of job-resume matching reveals that model scores are not attribute-neutral even when no demographic information is provided to a model [7]. The three mentioned studies clearly show that risks associated with biased filtering are not limited to traditional, rules-based systems but apply equally to modern LLMs that can potentially be used by a recruitment portal for resumes processing or questions generation.

B. Blind screening as a mitigating mechanism

Recruitment research conducted on auditing the talent pipelines of corporations by means of blind AI screening techniques contends that anonymizing identity-disclosing features (name, photo, gender, alma mater) prior to evaluation by a human or an algorithm diminishes the effect of the proxy factors on selection decisions and offers a methodology for auditing whether the pipeline is truly blind [2]. Another research on blind spots in recruitment via AI focuses on developing quantitative techniques for uncovering discrimination that persists even when identity attributes have been stripped out of resumes, for instance, discrimination based on proxies like institutional affiliation, location, or writing style [3]. Taken together, these two papers validate the use of blind screening as an effective but flawed technique, which must be used in combination with other measures in order to create a multi-layered system, as in the framework presented in this paper, which combines blind screening with independent auditing after the selection process.

C. Factors influencing biased selection decisions

However, bias in the selection process does not solely reside in the algorithm. The case study on federal character policy recruitment in Nigeria illustrates how biases such as political ethnicity and patronage system may affect merit-based selection in the presence of a formal rule-based recruitment process in the public domain [4]. In the context of campus placement, the case study provides an important lesson, in the sense that failure of fairness in selection can occur because of organizational and network biases apart from whatever kind of screening tool is employed. This is why a TPO facing fairness dashboard and continuous auditing is recommended in the proposed framework.

D. Challenges to the idea of algorithmic neutrality

A critical examination of AI-assisted talent selection questions the neutrality of algorithmic technology and points out that any assumption of neutrality could hide the bias that is already there, thereby making an analysis of hiring decisions less rigorous [5]. The paper guides the decision that has been made in this framework that each step of the pipeline needs to be examined for bias, including the test itself, and just because CGPA has been replaced by a test does not make it unbiased.

E. Measures of fairness and evaluation tools

A comprehensive review of fairness in AI-based hiring examines the issues, measures, and techniques involved in evaluating hiring algorithms, as well as future directions in research that can be undertaken, such as the importance of having an evaluation that includes multi-dimensional intersectional fairness and not merely a single measure [8]. This paper is highly relevant in guiding how evaluation is conducted in this research paper, which, unlike other researches that provide one aggregate pass rate, presents pass rate parity along six different demographic dimensions.

F. Quality of AI-assisted test items

Since the framework suggested makes use of the LLM in drafting competency test questions (Section IV), quality and validity of test items created by AI become relevant to the discussion. A field study on large scale on quality of AI-generated tests analyzes the performance of items created by AI versus those created by humans with regards to clarity, difficulty level, and content validity [9]. The present paper provides the foundation for the design choice that all questions generated using the LLM need human review-and-refinement (TPO) before being used (Section IV).

G. Existing campus recruitment systems

Two examples of recently developed systems depict the current level of automation in campus placement processes. Development of an intelligent campus placement portal is one such system that primarily aims at automation of registration, application processing, and interviews for both students and TPOs [10]. Similarly, another example of an AI-based college placement system also highlights the need for automation of processes and, in certain cases, resume-job matching [11]. In neither case, as per the reviewed literature, does the system implement any form of fairness process, competency-based screening in place of CGPA filter, or blind screening process.

TABLE I. COMPARISON OF EXISTING CAMPUS PLACEMENT SYSTEMS AND THE PROPOSED FRAMEWORK

Feature	Intelligent Placement Portal [10]	AI-Powered College Placement System [11]	Proposed Framework
Primary eligibility filter	CGPA-based	CGPA / resume-based	Competency test (CGPA-independent)
Application & interview workflow automation	Yes	Yes	Yes (retained)

Feature	Intelligent Placement Portal [10]	AI-Powered College Placement System [11]	Proposed Framework
Blind screening option	Not reported	Not reported	Yes — optional, TPO-controlled
Explicit fairness auditing	Not reported	Not reported	Yes — 5 checkpoints, 6 demographic dimensions
Effort-based bonus scoring	Not reported	Not reported	Yes — internships, projects, certifications
Validation reported	System description only	System description only	Simulation on N = 1,000 synthetic students

III. PROPOSED METHODOLOGY

Unlike the pilot testing of the framework at an institution, the evaluation through simulation uses a synthetic population of students, making the use of simulation an appropriate methodology in this instance for three reasons. The first one is that the framework is still being designed, meaning there are no actual implementation results. The second reason is that simulation enables the testing of the fairness properties on a big and varied synthetic population, avoiding any risk of harming an actual student if there are design flaws in the framework, since a placement decision is quite important for a student's life. Finally, since the ground truth attributes of the synthetic population (which include a hidden aptitude value never seen by the scoring mechanism) are known, it becomes possible to verify not only the fairness properties of the framework, but also its validity properties.

A. Synthetic Dataset Design

In the simulation, a population of N=1,000 artificial student records, each containing 21 attributes, is considered. The academic status of the candidates is characterized by CGPA attribute whose values are spread within the realistic range of 3.5 – 9.9, where students are separated into "Average-CGPA" category (which consists of those who have the scores below the traditional 7.0 eligibility threshold) and "Higher-CGPA" category (with those who have the scores equal to or greater than it); 407 students (40.7% of 1,000 total number of students in this population) belong to Average-CGPA category. There is gender parity (with 498 female students). All the students receive the Prior Domain Exposure score (Low / Medium / High) which stands for the hypothetical internships, coaching and industry experience and Latent Aptitude score which stands for the true competency of the student (job-related) and is utilized solely for test results generation.

Nine more TPO-provided variables were included in the model to emulate a data set reflecting the effort-based information that a placement office can reasonably gather independently of the candidate's background: prior interview experience, internship experience count, project experience, certification experience count, self-rated communication skill (scale from 1 to 5), technical skills diversity, and, solely for auditing purpose, socioeconomic background code, origin region code, and first-generation student. In line with the design principle of the framework described in Section IV, the socioeconomic and regional variables are employed only for the audit performed in Section V and never serve as scoring or filtering inputs.

B. Test-Score Generation

The student-specific simulated performance on the competency test (Test_Score_100, transformed to Competency_Score_0to1) is created as a function of Latent Aptitude and Prior Domain Exposure, plus noise to capture day-to-day variation and imperfection in measurement. The simulation explicitly introduces a ground-truth relationship between the two variables to ensure that the exposure-pass rate gradient described in Section V (Table/Chart 6) can be used as a validity check; if it does not appear in the simulated test results, then the test clearly does not measure competence.

C. Comparison Mechanisms

Two methods of selection are applied to the same population and are contrasted against each other. The OLD system is a continuation of tradition: those whose hard CGPA is above 7.0 are selected to be considered by the company, while nothing else is taken into account.

The NEW method is an application of the four-stage process described in Section IV: everybody takes a competency test, and the criterion of selection at the competency stage is a low, deliberately permissive pass threshold (the score above 20% of the possible maximum), so that this stage can act as a minimum-competency gate, not a competitive one; the candidates who pass the competency gate are optionally anonymous through blind screening; and then the top 30% among those who have passed are selected and sent to the recruiting company as the shortlist. This system of double gating allows separating the questions “deserves a candidate to have an evaluation opportunity” and “who is the best candidate?” which are mixed together in the single CGPA-only baseline.

D. Fairness Metrics

The fairness metric is computed based on the pass rates parity between six demographics — gender, CGPA category, geographic location, socio-economic background stand-in, first-generation, and previous domain exposure based on the recommendation from the fairness-metric literature to use a multi-dimensional evaluation approach [8]. A pass rate for any particular subgroup is deemed to satisfy the fairness criterion if it is above 90% (the minimum inclusion threshold in prior fairness assessments) and if it approaches demographic parity, which requires it to be above 95%.

The treatment of the dimension of previous domain exposure is different from those of the other five dimensions, in that in exposure, an increasing gradient of pass rates is expected and desirable, since exposure is expected to correlate with competency, and otherwise, the test will not have measured anything meaningful.

IV. IMPLEMENTATION

The proposed architecture is structured as a four-step pipeline, where fairness testing has been designed at each stage, depicted in Fig. 1 below. This section explains the system architecture and module design; as mentioned in Section I, the framework is still in design stage, and not in operation, hence the explanation provided refers to the designed system and not the existing software.

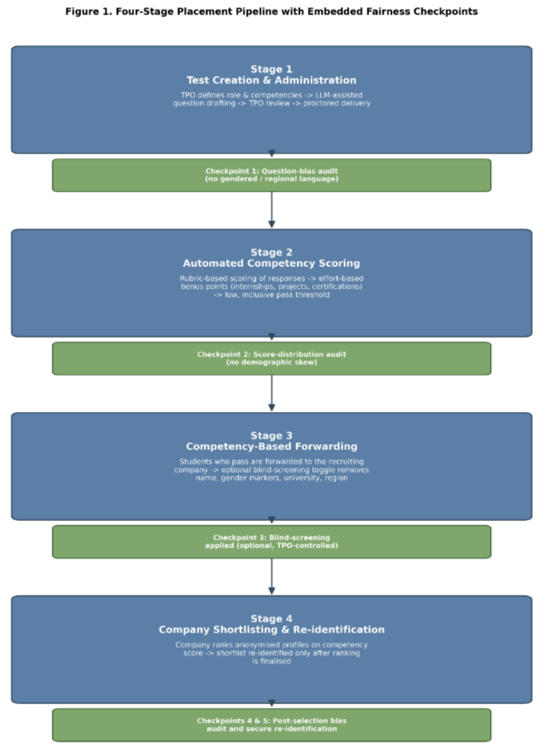

Fig. 1. Four-stage placement pipeline with five embedded fairness checkpoints.

A. Core Modules

- TPO Dashboard: the control interface where a Training and Placement Officer creates a role, determines its competencies, and configures its pipeline settings (whether the pipeline will employ blind screening or not). It also displays fairness audit reports for each cycle.
- Test Question Creation Module: uses an LLM to create draft questions on candidate competencies, which are then edited by the TPO. This process is driven by the need for human review of test questions created by AI [9].
- Scoring and Ranking Module: implements rubric scoring on test results; provides effort-based bonuses (for number of internships, projects and certificates); uses low inclusion threshold for competencies as described in Section III-C.
- Blind Screening Module: optional, TPO-controlled module that removes name, gender-related pronouns, university affiliation and regional identifiers from the candidate’s profile before it is presented to the recruiting firm, with re-identification done only after the ranking of the recruiting firm is finalized based on corporate blind

screening procedure [2]; specifically constructed to address proxy variable leakage problem [3].

- Fairness Auditing Module: Calculates the six-dimension pass rate audit described in Section III-D after each placement run, and notifies about any subgroup whose pass rate is below the required fairness threshold, implementing the functionality of Checkpoints 1, 2, 4, and 5 in Figure 1.

B. Technology Rationale

For any future implementation, a locally hosted LLM, open weight (an example of which can be found within the Llama family of LLMs), will be used for question generation, rather than an external API hosted solution, mostly due to data privacy and cost predictability considerations, as no student assessment data would have to exit the institutional network. The backend system will consist of a relational database and Django framework, among other technologies, selected due to its maturity and administrative capabilities.

V. RESULTS AND DISCUSSION

A. Simulation Setup

Both the OLD (CGPA ≥ 7.0) and NEW (Competency-Pipeline) systems, discussed in Section III-C, were employed on the artificial population (N = 1,000; Section III-A), and the results were collected at every pipeline stage for each of the students.

B. Headline Result: Opportunity Creation

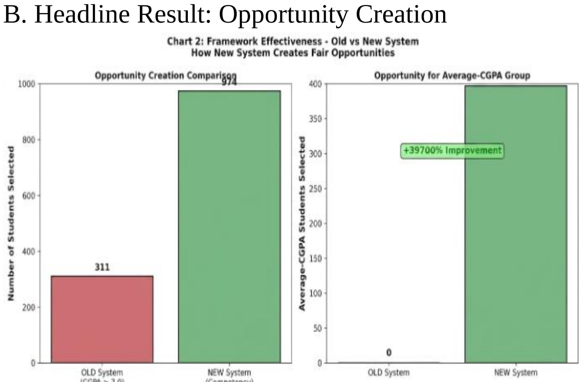

Fig. 2. Opportunity creation: the OLD (CGPA-only) system advances 311 of 1,000 students (31.1%) to assessment, versus 974 of 1,000 (97.4%) under the NEW competency-gate system; among Average-CGPA students specifically, representation rises from 0 selected under the OLD system to 397 (97.5% of that subgroup) under the NEW system.

Under the OLD approach, 311 out of 1,000 students (31.1%) pass the criterion of CGPA ≥ 7.0 and thus are even considered in the first place; the other 689 students (68.9%) do not make the cut and are disqualified prior to assessing any skill. Under the NEW approach, 974 out of 1,000 students (97.4%) pass the relatively low, more inclusive competency barrier and reach the pool of assessors. The most illustrative example can be given by the Average-CGPA group, where among the 407 students who are below the conventional barrier of 7.0, none will be selected under the OLD approach but 397 (97.5%) will pass the competency gate under the NEW approach. This result is in agreement with the design of the new framework, which was intended to separate the competency screening process from the CGPA [1], [5].

C. Validity Check: Exposure-Based Pass-Rate Gradient

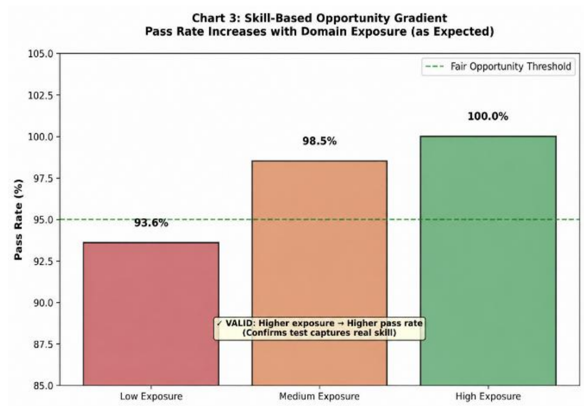

Fig. 3. Competency-test pass rate by prior domain exposure: 93.6% (Low), 98.5% (Medium), 100.0% (High).

Since a lenient pass criterion could, in principle, result in passing almost everyone irrespective of their skills, an internal validity test should be conducted before accepting the high pass rate observed in Section V-B as proof of a fair test and not as a mere by-product of the testing methodology used. The pass rate increases in a monotonic fashion with increasing level of domain exposure, ranging from 93.6% for Low-exposure participants, 98.5% for Medium, and 100.0% for High. This is consistent with what one would expect if the test was actually measuring the competency signal that has been embedded in the simulation based on exposure levels (Section III-B).

D. Multi-Dimensional Fairness Audit

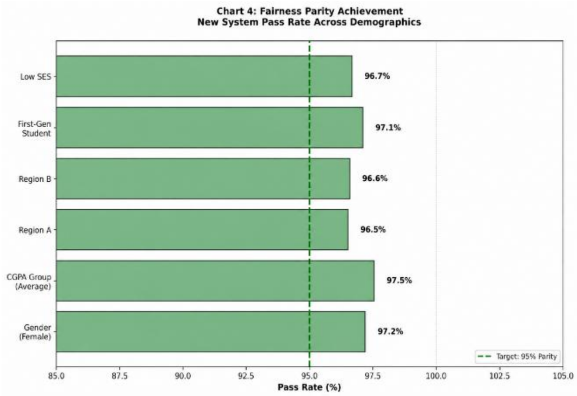

Fig. 4. Competency-gate pass rate across six demographic dimensions, all above the 95% target line.

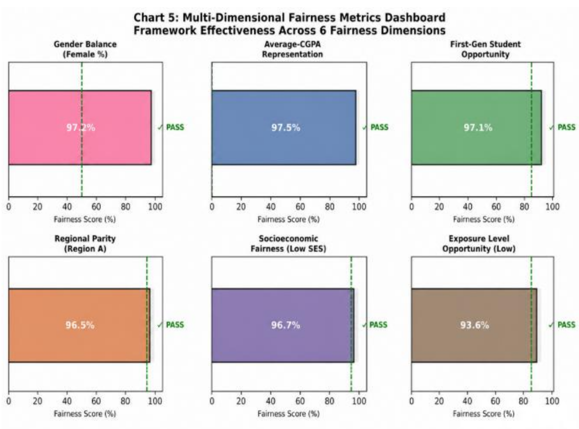

Fig. 5. Six-dimensional fairness dashboard; every audited dimension passes its fairness threshold.

In accordance with multi-dimensional assessments in the literature on fairness [8], the pass rate audit was performed separately for each of the six subgroups rather than being provided as one consolidated statistic.

All subgroups analyzed (females at 97.2%, the Average-CGPA group at 97.5%, Region A at 96.5%, Region B at 96.6%, first-generation students at 97.1%, and the low socio-economic proxy group at 96.7%) meet both the floor threshold and the parity threshold established in Section III-D.

None of the subgroups have pass rates significantly lower than the overall population, which is the desired result to achieve in order to mitigate the demographic leakage observed in prior recruitment-bias studies [1], [3], [6], [7].

E. Pipeline Retention

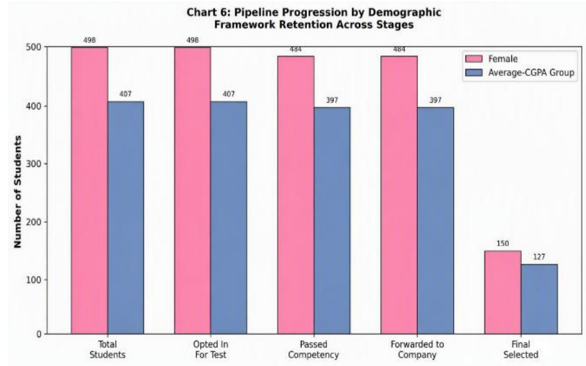

Fig. 6. Stage-by-stage retention of Female and Average-CGPA students through the four-stage pipeline; both subgroups remain proportionally represented at every stage, including the final top-30% company shortlist.

Parity between the two subgroups at the same level in a particular pipeline stage may hide the attrition process occurring over stages gradually. The follow-up of the female and average CGPA subgroups from the very first stage to the last one, i.e., from opting for the competency test to the company shortlist based on top 30%, reveals the same retention rate among the two subgroups at each and every stage: 484 out of 498 female students (97.2%), and 397 out of 407 average CGPA students (97.5%) get past the competency test gate; also proportional representation persists further in the forward stage to reach the top 30% list (150 out of 498 female, and 127 out of 407 average CGPA students selected).

F. Discussion

These findings largely comport with the issues discussed in the literature analyzed in Section II and provide simulations demonstrating the value of combining, as opposed to selecting, the mitigation techniques described in that literature. The opportunity creation finding (Section V-B) is a response to the criticism regarding the use of proxy data in [1] and [5]: replacing a filter that is correlated with background information with a low threshold competency gate significantly alters the set of individuals who are even considered. The multidimensional audit (Section V-D) represents the approach suggested by [8] in lieu of using a one-dimensional measure of fairness which prior research warns can mask differential treatment in subgroups. The validation test conducted in Section V-

C responds to an issue implied by [3] and [5]: a fairness mechanism that approves all applicants is not fair, just uninformative, and the exposure gradient finding is stated to preclude such a conclusion.

G. Limitations

The findings must be taken as proof-of-concept evidence rather than evidence that the model works in a practical setting. Four key limitations apply. Firstly, all the data used is generated: the correlation between Latent Aptitude, Prior Domain Exposure and Test_Score_100 is a modeling assumption made by the researchers, and not a relationship that was empirically measured, and the distribution of test scores might well behave differently in practice. Secondly, this simulation models a generic position only; the requirements for competency and hence the dynamics of unfairness will vary greatly depending on role (for example, core engineering position versus analyst versus management trainee). Thirdly, the simulation does not take into account the recruiter's behavior at the final interview stage, when human biases discussed in [1], [4] and [6] can still emerge despite the generation of a fair shortlist – blind screening in this model is limited to the shortlist generation stage, not to interviews. Finally, no real students, TPOs and recruiters were included in the simulation – real-world testing using institutional data and, eventually, live pilot programs at two or three institutions would be necessary before recommending implementation.

VI. CONCLUSION AND FUTURE SCOPE

The CGPA filter alone filters out many candidates from placement considerations on campus even before assessing them for any competencies related to job performance in a manner that is aligned with the general trends observed about biased practices of selection procedures both algorithmic and traditional. In this paper, we propose a four-tier process for placement of students which replaces the strict CGPA filter with a competency test, effort-based score calculation, and optionally blind screening of candidates, and tested our proposal by simulating it over a synthetic population of 1,000 students. Our simulation showed a 97.4% rate of assessment opportunity for candidates as opposed to 31.1% under CGPA baseline, and 93%+ pass-rate equality across six demographic axes. These results must be seen as concept validation via simulation instead of as evidence of real-world efficacy. The following steps will take this research closer to implementation: (1) a pilot study with two to three partner organizations using actual student data and one well-defined role to validate whether the effort-based bonus score system and the competency gate operate in reality as they did in the simulations; (2) an expansion of the current framework to include additional, structurally distinct roles, because the fairness dynamics are not likely to be the same in a software engineering role versus a core engineering role; and (3) a bias audit from the perspective of the recruiter at the interview stage to see if the advantages of the blind screening disappear when the candidates are again identified for interviews.

ACKNOWLEDGMENT

Acknowledgements: The authors are thankful to Dr. Sunil Mahajan Associate professor for his invaluable guidance in this study, and also to the MSC Information Technology of MIT Arts, Commerce and Science College, Alandi (D), Pune, Maharashtra, India.

REFERENCES

[1] El Moutafail and K. Belkhoutout, "Bias and fairness in LLM-based recruitment: A systematic review," *EPJ Web of Conferences*, vol. 381, Art. no. 00015, 2026, AICSET 2026.
[2] T. Baron and M. Milton, "Eliminating hiring bias in corporate recruitment: Auditing talent pipelines using blind AI screening tools," Avance International Journal of Business and Finance Research, vol. 2, no. 7, pp. 208–221, Jul. 2026.
[3] A. I. Novitskaya, "AI blind spots in recruitment, hidden discrimination, and methods for its quantitative detection," International Journal of Open Information Technologies, vol. 14, no. 8, pp. 140–145, Aug. 2026.
[4] O. O. Godwin and A. Joel, "Federal character recruitment in Nigeria: Ethnic politics, patronage networks and the crisis of meritocracy (2020–2026)," Port Harcourt Journal of Society and Environment, vol. 4, no. 1, pp. 257–271, Jun. 2026.

[5] R. Herrera Martínez, Y. Mucharraz Y Cano, and K. Cuilty, “The myth of neutrality in AI-based talent selection,” Academy of Management Proceedings, 2026, doi: 10.5465/AMPROC.2026.18508abstract.

[6] A. K. Veldanda, F. Grob, S. Thakur, H. Pearce, B. Tan, R. Karri, and S. Garg, “Investigating hiring bias in large language models,” in R0-FoMo: Workshop on Robustness of Few-Shot and Zero-Shot Learning in Foundation Models, NeurIPS, 2023.

[7] H. Iso, P. Pezeshkpour, N. Bhutani, and E. Hruschka, “Evaluating bias in LLMs for job-resume matching: Gender, race, and education,” in Proc. 2025 Conf. North American Chapter Assoc. Comput. Linguistics: Human Language Technologies (Industry Track), 2025, pp. 672–683.

[8] D. F. Mujtaba and N. R. Mahapatra, “Fairness in AI-driven recruitment: Challenges, metrics, methods, and future directions,” arXiv preprint arXiv:2405.19699, 2024, revised 2025.

[9] C. Isley, J. Gilbert, E. Kassos, M. Kocher, A. Nie, E. Brunskill, B. Domingue, J. Hofman, J. Legewie, T. Svoronos, C. Tuminelli, and S. Goel, “Assessing the quality of AI-generated exams: A large-scale field study,” in Proc. AAAI Conf. Artificial Intelligence (AAAI-26), 2026.

[10] A. Rahangdale, O. Deshmukh, R. Biradar, Y. Gade, and M. Powar, “Design and development of an intelligent campus placement portal,” Int. J. Engineering Research & Technology (IJERT), vol. 15, no. 6, Jun. 2026.

[11] A. Yadav, A. K. Saw, and A. K. Gupta, “AI-powered college placement system,” *Int. Research J. Modernization in Engineering Technology and Science (IRJMETS)*, vol. 8, no. 2, pp. 2249–2253, Feb. 2026.