A Readiness-Aware Internship Recommendation and Skill-Gap Analysis System for Students

Eshan Metwate[1], Sairaj Patil[2], Ayush Toke[3], Lakhan Hirave[4], Dr. Sunil Mahajan[5]
[1,2,3,4]*Student, Department of Computer Science, MIT Arts, Commerce and Science College, Alandi (D), Pune, Maharashtra, India*
[5]*Associate Professor, Department of Electronics, MIT Arts, Commerce and Science College, Alandi (D), Pune, Maharashtra, India*
doi.org/10.64643/ijirt.208969-459

***Abstract*—Students today have access to so many internship listings than ever, spread across various internship-search platforms, college placement portals, company career pages, government-run schemes. Yet most of these platforms focus primarily on matching and ranking opportunities, while providing limited support for understanding whether a student is actually eligible, which specific requirements they do not satisfy, and what skills they could improve. This paper proposes a conceptual architecture for a readiness-aware internship recommendation and skill-gap analysis system for students. The proposed design compares a student's education, skills, experience, and preferences against internship requirements and extends conventional matching by identifying satisfied and unsatisfied eligibility conditions and highlighting specific skill gaps. Rather than presenting a system that has already been implemented and empirically validated, this paper defines its conceptual architecture, data flow, processing stages, and reasoning logic as a theoretical foundation for future implementation and empirical evaluation. The proposed framework is situated within gaps identified across internship recommendation, explainable matching, and skill-gap analysis research. The architecture therefore provides a basis for future investigation into whether presenting readiness and skill-gap information alongside internship recommendations can help students better understand requirements and make more informed application decisions.**

***Index Terms*—Readiness-Aware Internship Recommendation, Internship Readiness, Skill-Gap Analysis, Student Decision Support, Student Employability, Career Decision Support**

I. INTRODUCTION

An internship is often the first point at which a student's academic preparation is tested against the practical expectations of an employer, and for many students it functions as a decisive early step toward employment. Over the past several years, the number of internship postings visible to students has expanded considerably, spread across college placement portals, corporate career pages, government-run internship schemes, and dedicated aggregator platforms. This expansion has undeniably widened access to opportunities, but it has also produced a familiar side effect: an abundance of listings that no student can realistically evaluate one at a time. Faced with postings that differ in eligibility rules, required skills, academic-year restrictions, and preferred prior experience, students are largely left to work out on their own which of these opportunities they are actually in a position to pursue.

Research on this problem has, for the most part, treated it as a matching or ranking task. Systematic reviews of job and internship recommender systems describe an established toolkit of techniques content-based filtering that compares resume text against posting text, collaborative filtering that draws on patterns across similar users, and hybrid combinations of the two all directed at producing an ordered list of postings judged most relevant to a given profile [1], [2]. A broader challenge-based survey of e-recruitment recommendation systems frames the domain in similar terms, noting that because e-recruitment carries lasting career consequences and because job seekers typically interact with a platform only briefly before leaving it once employed, standard recommendation techniques tend to perform less reliably here than in domains such as retail or entertainment [3]. What these reviews share is a largely implicit assumption: that the value of a recommender system lies chiefly in surfacing the right opportunities, with the separate question of whether a student is actually ready for a surfaced opportunity left for the student to work out independently.

A more recent strand of work has begun to question this assumption from a different angle, arguing that as person-job recommendation systems increasingly rely on deep and often opaque models, their outputs risk becoming difficult for either party to interpret. A systematic review of explainable person-job recommendation systems catalogues techniques feature attribution, attention mechanisms, knowledge-graph reasoning, counterfactual explanation developed specifically to open up this black box [4]. Much of this literature, however, is oriented toward the hiring side of the transaction, toward helping recruiters trust a ranking or detect bias within it, rather than toward helping an individual applicant understand, in concrete terms, why they were or were not judged a strong match for a specific role, and what would need to change for that judgment to shift.

A separate, largely parallel body of work addresses skill gaps and employability directly. Some of this research builds systems that extract skills from student resumes and compare them against what internships or job postings request, explicitly aiming to narrow the distance between what students learn in the classroom and what employers expect at the point of hiring [5]. Other studies model student readiness statistically, training classifiers on academic and behavioural data to estimate how likely a student is to be placed. These efforts confirm that a student's readiness for the job market can be estimated with reasonable accuracy, but they are generally built and evaluated as institutional analytics tools, sitting apart from the recommendation process itself rather than being attached to individual, real-time internship suggestions.

Taken together, these three strands of research recommendation and ranking, explainability, and skill-gap or employability prediction rarely intersect. Existing internship recommendation approaches primarily emphasise opportunity matching and ranking, while comparatively little attention has been paid to evaluating a student's readiness for a specific internship and to presenting whatever is missing in a form the student can act on directly, at the point an application decision is being made. This raises a question that, to the best of our knowledge, has not been directly examined in the internship context: does presenting readiness and skill-gap information alongside a recommendation actually improve a student's decision-making, relative to the conventional approach of searching and applying from a ranked list alone?

This paper does not attempt to answer that question empirically. It takes the necessary first step toward doing so: it proposes a formal conceptual architecture for a readiness-aware internship recommendation and skill-gap analysis system, one that compares a student's education, skills, experience, and preferences against internship requirements and is designed, at a structural level, to surface not just a match score but the specific eligibility conditions a student satisfies, the ones they do not, and the skills that would close that distance. The architecture presented here is intended to serve as the theoretical basis for subsequent implementation and empirical evaluation, rather than as a report of a system that has already been built and tested.

The remainder of this paper proceeds as follows. Section II reviews the literature underpinning this proposal across five related areas internship-specific recommendation, broader job-recommender-system surveys, explainability in recommendation, skill-gap and employability assessment, and career guidance and decision-support systems before closing with a synthesis of the gap this paper addresses. Subsequent sections, to follow in later stages of this work, will set out the proposed system's architecture and the methodology by which its conceptual claims can eventually be tested empirically.

II. LITERATURE REVIEW

This section situates the proposed system within five overlapping strands of existing research: internship-specific recommendation systems, broader systematic reviews of job recommender systems, work on explainability in recommendation, research on skill-gap analysis and employability prediction, and career guidance and decision-support systems more generally. It closes by drawing these strands together to state precisely the gap this paper addresses.

A. Internship-Specific Recommendation Systems

Several studies have targeted internship recommendation specifically, rather than job recommendation in general. One of the earliest used an artificial neural network trained on competence-test and questionnaire data to recommend internship placements to informatics students, reporting near-perfect recognition on training data and a high accuracy rate on previously unseen cases, although the reasoning behind any individual recommendation remains internal to the trained network and is not exposed to the student [6]. A systematic review by Poncio surveyed the internship-profile-matching literature specifically, examining how researchers collect the data used to build a student profile and which classification algorithms and ranking metrics are used to generate suggestions; the review reports that data collected from primary sources tends to yield more personalised results than data drawn from secondary sources, and that hybrid recommenders are generally used to offset the limitations of either approach on its own [1]. More recent, applied systems have moved toward semantic matching: one internship recommendation engine built for a large government internship scheme uses sentence-transformer embeddings and vector-similarity search to match a student's skills, field of study, and location preferences against postings, paired with a skill-gap-analysis module intended to indicate what separates a given student from a stronger applicant [7]. A related, Java-based system takes a different route to a similar end, combining rule-based filtering with vectorised similarity scoring across skills, academic attributes, and stated preferences, and deliberately avoids machine-learning models whose internal reasoning cannot be inspected, so that the logic behind each recommendation stays legible to students and administrators alike [8]. A broader summary of AI-based internship recommendation work reaches a comparable conclusion across these systems: their central contribution has been reducing information overload and improving profile-to-posting alignment, not assessing whether a student meets an internship's specific requirements [9].

B. Job Recommender Systems: Broader Systematic Reviews

Looking beyond internships to job recommendation more broadly, two recent reviews offer a wider view of where the field currently stands. A systematic literature review covering roughly thirteen years of research on job recommender systems catalogues the algorithmic approaches used throughout this period and the metrics by which they are typically evaluated, and observes that despite the field's steady growth, no comprehensive review had previously synthesised it in this way, a gap the review sets out to close [2]. Rather than organising the field by algorithm family, a challenge-based survey of e-recruitment recommendation systems instead frames the domain through the practical difficulties it poses for developers: the high stakes attached to a poor recommendation, the short and sparse interaction histories typical of job seekers, and the frequent turnover of postings, all of which distinguish e-recruitment from more conventional recommendation domains [3]. Work on embedding-based matching has continued to refine this ranking-centred paradigm; one recent model embeds resumes and a standardised occupational taxonomy into a shared vector space specifically to improve the generality of job matching across roles, again without addressing whether a given applicant satisfies that role's stated requirements [10]. Neither of the two reviews, nor the embedding-based work that follows in their tradition, treats applicant-side readiness assessment as a first-class design concern; all are organised, implicitly or explicitly, around the goal of producing an accurate ranked list.

C. Explainability in Recommendation Systems

A separate line of research questions whether an accurate ranking is, on its own, sufficient. As person-job recommendation systems have adopted increasingly complex deep-learning architectures, a systematic review of explainable person-job recommendation methods argues that many of these architectures function as black boxes, producing a ranked output without a legible account of why that ranking was produced. The review organises explainability techniques into three layers the data used, the model's internal reasoning, and the final output and reports that counterfactual explanation methods, which describe what would need to change about a candidate for an outcome to differ, achieved the strongest balance between explanatory power and predictive performance among the methods it compared [4]. This finding is directly relevant to the problem addressed here, since a counterfactual explanation is, in structure, close to a gap statement: what is currently missing, and what would close it. The review's focus, however, remains on hiring platforms and on the trust of recruiters and job seekers in aggregate, rather than on equipping an individual student with an explanation they can use to decide whether, and how, to apply for a specific opportunity.

D. Skill-Gap Analysis and Employability Prediction
A further strand of research treats skill gaps and employability as problems to be modelled directly, generally at some remove from the recommendation process itself. One system pairs a resume-parsing layer with a component that identifies missing skills and suggests specific areas for improvement, framing its objective explicitly as narrowing the distance between what students are taught and what employers require [5]. Elsewhere, researchers have trained supervised classifiers on academic and behavioural data to predict employability, with one comparative study of k-nearest neighbours, random forests, naive Bayes, and support vector machines finding that the k-nearest-neighbours classifier outperformed the other three on this task [11]. A related system, built specifically to output a placement-readiness judgment for individual students, combines similar academic and behavioural indicators to produce both a binary readiness outcome and an accompanying confidence score [12]. Taken together, these employability-prediction studies show that readiness can be estimated with reasonable statistical confidence; what they generally do not do is attach that estimate, and the reasoning behind it, to a specific internship recommendation at the moment a student is deciding whether to apply.

E. Career Guidance and Decision-Support Systems
Finally, a body of work outside the internship-recommendation literature narrowly defined addresses career decision-making using similar tools at a broader scope. Some recent systems use large language models to conduct a structured, multi-step assessment of a student's academic background, interests, and goals before generating personalised career recommendations [13]. Others rely on more classical, interpretable classifiers decision trees, random forests, and naive Bayes paired with a conversational interface, to guide students toward a suitable stream or career path after school [14].
These systems demonstrate that structured, explainable reasoning can be layered on top of a recommendation engine in a way students find usable, but their scope is typically the broad question of which career or educational stream to pursue, rather than the narrower and more immediate question of whether a student is ready to apply for one specific internship listing today.

F. Summary and Research Gap
Drawing these strands together, a consistent pattern emerges. Internship-specific recommendation systems and the broader job-recommender-system literature are both organised principally around producing an accurate ranked list, whether through classical filtering techniques, semantic embeddings, or newer taxonomy-aware models [1], [2], [3], [6], [7], [8], [9], [10]. Explainability research has begun to open up the reasoning behind such rankings, and counterfactual methods in particular already resemble, in structure, a gap statement, yet this work remains oriented toward hiring platforms and aggregate trust rather than toward an individual student's application decision [4]. Skill-gap and employability-prediction research shows that readiness can be estimated with reasonable accuracy, but almost always as a standalone institutional or diagnostic tool, disconnected from any specific recommendation [5], [11], [12].

Career-guidance systems bring interpretable, student-facing reasoning into view, but at the scale of an entire career or academic stream rather than a single internship listing [13], [14]. The reviewed literature provides limited evidence of systems that integrate internship recommendation. This is the gap the present paper addresses. It proposes a formal conceptual architecture for a readiness-aware internship recommendation and skill-gap analysis system, intended to serve as the theoretical foundation for the implementation and empirical evaluation needed to test whether readiness and skill-gap information, presented alongside a recommendation, genuinely improves students' internship decisions.

III. PROPOSED METHODOLOGY / SYSTEM DESIGN

The architecture proposed here responds directly to the gap set out in Section II: existing internship-recommendation systems tell a student what is available, but not whether they are ready for it, or what stands in the way if they are not. What follows is organised as five conceptually distinct layers, moving from raw student and internship data through to a ranked, explained recommendation. Consistent with the scope of this paper, this is a structural proposal, the reasoning behind each layer and how it connects to the next rather than a description of a system that has already been coded and deployed.

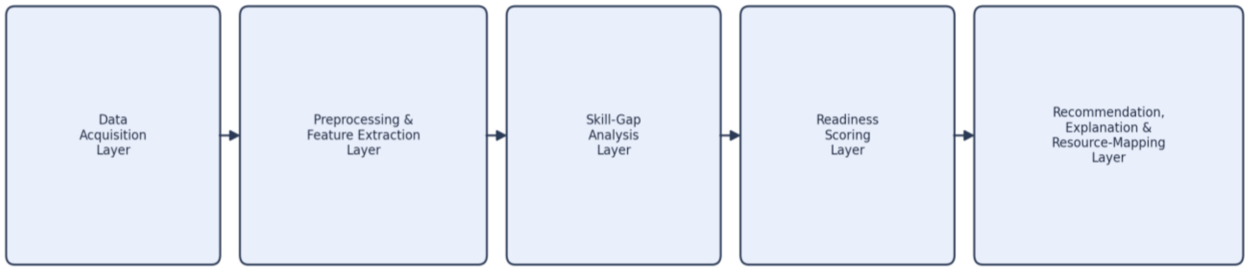

Figure 1. Proposed five-layer architecture for the readiness-aware internship recommendation and skill-gap analysis system. Data flows left to right; each layer passes a structured output to the next

A. Data Acquisition Layer

This layer is where the two inputs the rest of the architecture depends on would enter the system: a student's profile academic record, self-reported skills, prior projects, internships and an internship posting's stated requirements required and preferred skills, eligibility rules, and role description. Treating data acquisition as its own layer, rather than as an incidental first step, is a deliberate choice: every layer after it assumes both inputs already exist in a comparable, structured form, and getting that form right is what makes the rest of the pipeline possible.

B. Preprocessing and Feature Extraction Layer

Student resumes and internship postings are rarely written in comparable language, so this layer is responsible for turning free text resume bullet points, project descriptions, posting requirements into a standard set of skill terms that both sides of the comparison can be measured against. Formatting artefacts, filler phrases, and inconsistent capitalisation would be stripped at this stage, leaving a clean, tokenised list of candidate skills, tools, and technologies for the layer that follows.

C. Skill-Gap Analysis Layer

With both profiles reduced to a common skill vocabulary, this layer compares them directly: for a given student and a given internship, it works out which required skills the student already has, which they partially have, and which are missing altogether. A practical case illustrates why this comparison needs to look past exact wording. An internship listing might ask for experience with the "MERN stack," while a student's resume lists React, Node, Express, and MongoDB individually, without ever using the umbrella term. A system built around literal string overlap would miss this match; the architecture is designed instead to compare skills by what they mean, using semantic similarity rather than keyword matching, so that equivalent ways of describing the same competency are recognised as such [10]. This is also what allows the layer to report a genuine gap a skill the student truly lacks rather than a false one caused only by a difference in terminology.

D. Readiness Scoring Layer

The output of the previous layer matched, partially matched, and missing skills is aggregated here into a single readiness figure for a given student–internship pair, informed by approaches used elsewhere to model student employability from structured indicators [11], [14]. Not every requirement is proposed to count equally: a posting's specific, high-signal requirements (a named framework or platform) would be weighted more heavily than generic, low-information phrases a resume might use to describe a candidate a broad self-description such as "strong communicator," for instance. This weighting is intended to keep the readiness score meaningful even when a resume leans on generic language, a known weakness of simpler, count-based matching approaches.

E. Recommendation, Explanation and Resource-Mapping Layer

The final layer is where a readiness score becomes something a student can act on. Internship opportunities would be ranked using the readiness score from Layer D, and each recommendation would be accompanied by a plain-language explanation of what qualified the student for it and, where relevant, what is still missing in the spirit of the counterfactual, gap-style explanations shown elsewhere to aid interpretability in person-job recommendation [4]. Where a specific skill gap is identified, this layer is also designed to attach it to a relevant external learning resource, so a student is pointed toward closing the gap rather than simply told one exists. Exactly how this output would be surfaced to a student a dashboard, a ranked list, or another format is an interface question left to the implementation stage in Section IV; what this layer defines is the information that interface would need to display.

F. Dataset and Tooling Considerations

No dataset has been collected or used at this stage of the work. Conceptually, realising the layers above would require two categories of data: a set of internship postings with clearly stated skill requirements, and a set of student profiles containing self-reported or verified skill information. A shared skill taxonomy a common vocabulary both sides are mapped against is proposed as a prerequisite for reliable comparison in Layer C, since without one, differently worded postings and profiles are difficult to compare consistently [10].

At this conceptual stage, "tools" refers to the categories of technique the design assumes rather than a fixed software stack, which Section IV addresses separately.

The architecture assumes natural-language processing methods for the extraction work in Layer B, embedding-based similarity search for the comparison in Layer C, and a lightweight scoring or classification approach for Layer D categories chosen because comparable combinations have already been explored in adjacent internship and career-recommendation research [5], [7].

IV. IMPLEMENTATION

Section III defines what the system is meant to do; this section outlines, at a planning level, what building it would actually involve. Nothing described below has been implemented the technology choices and module breakdown here are a proposed path from the architecture in Section III to a working prototype, not a record of one that already exists.

A. Proposed Technology Stack

A conventional web-application stack is proposed, consistent with similar academic career-guidance and internship-matching prototypes described in the literature [8], [9], [13]. The intended stack is summarised in Table I.

TABLE I. PROPOSED TECHNOLOGY STACK

Layer	Component	Purpose
Backend	Python (Flask / FastAPI)	Application logic and API layer
Skill Extraction & Matching	spaCy / NLTK + Sentence-Transformers (SBERT)	Free-text parsing and embedding-based similarity
Readiness Scoring	scikit-learn	Implementing and evaluating classification models
Data Storage	MongoDB (or a relational alternative)	Storing student profiles, postings, and computed readiness scores
Frontend	React.js	Student-facing dashboard and interface

B. Setup Overview

At an implementation stage, the system would be structured as a set of independently deployable services corresponding to the five layers in Figure 1, so that the preprocessing, scoring, and recommendation components could be updated without affecting one another.

This mirrors the modular deployment pattern used in comparable AI-based recommendation and career-counselling tools [13], and follows directly from the layered design rationale set out in Section III.A.

C. Module-Wise Description

Table II maps each architectural layer from Section III.A onto its intended function and technology category at implementation stage; this mapping is a planning reference for future work, not a record of components already built.

TABLE II. MODULE-WISE DESCRIPTION

Module	Intended Function	Technology Category
Data Acquisition	Collects student and internship posting data	Web forms / API ingestion
Preprocessing & Feature Extraction	Converts free text into structured skill data	NLP pipeline
Skill-Gap Analysis	Compares student and posting skill sets	Embedding-based similarity
Readiness Scoring	Produces an interpretable readiness score	Classification model
Recommendation & Explanation	Ranks opportunities and explains each result	Rule-based ranking, rendered through the dashboard in Section III.E

V. RESULTS AND DISCUSSION / PROPOSED EVALUATION FRAMEWORK

No prototype exists yet, so there are no results to report in the usual sense. What this section lays out instead is the evaluation plan this architecture is meant to be tested against once it is built the metrics, the comparison point, and the outcomes the design in Section III would predict.

A. Proposed Evaluation Metrics

Table III sets out the six measures against which the proposed architecture is intended to be evaluated, consistent with the criteria stated in the abstract.

TABLE III. PROPOSED EVALUATION METRICS

Metric	Purpose
Matching accuracy	Whether recommended internships are genuinely relevant to a student's profile, against a keyword-based baseline
Decision-making time	How long students take to decide whether to apply, against conventional manual searching
Understanding of requirements	Whether students can correctly state what an internship requires after viewing a recommendation
Skill-gap identification accuracy	Whether the gaps the system flags match those a subject-matter expert would identify
Usability	How easily students can interpret and act on the dashboard output described in Section III.E
Confidence in application decisions	Whether students feel more certain about whether to apply after seeing readiness and skill-gap information

These measures follow evaluation patterns already used to assess student employability-prediction models and explainable recommendation systems, which similarly combine predictive accuracy with human-judgment and usability measures rather than relying on ranking accuracy alone [4], [11], [12].

B. Proposed Comparative Baseline

A natural baseline for future evaluation is a conventional keyword- or requirement-matching recommender, of the kind critiqued in the systematic reviews discussed in Section II [1], [2], [3]. Comparing the proposed readiness-aware approach against such a baseline would help isolate whether the added skill-gap and readiness layers meaningfully change recommendation quality and decision-making, rather than simply reproducing what a simpler matching system would already return.

C. Illustrative Walkthrough

To make the intended behaviour of the readiness and skill-gap outputs concrete, consider a hypothetical student profile built around web-development coursework and projects, matched against a typical web-development internship listing. Where the internship's stated requirements include REST API integration and version control with Git, and the student's profile does not mention either explicitly, the system is designed to surface both as explicit skill gaps rather than folding them into a lowered overall score, and to attach each gap to a corresponding learning resource through the resource-mapping component described in Section III.E. This example illustrates the system's intended logic; it is not a reported finding.

D. Anticipated Outcomes and Discussion

Based on the design rationale in Section III, the proposed system is expected to differ from conventional recommenders in two main ways. First, by surfacing partial matches and explicit skill gaps rather than only strict matches, it should reduce cases where a student either overlooks a suitable-but-imperfect opportunity or applies to one they are clearly underprepared for. Second, by attaching an explanation to each recommendation, it is expected to improve student trust and comprehension relative to opaque, ranking-only systems, consistent with the concerns about transparency raised in Section II.C [4]. These are stated as expected directions to be validated once a working prototype and a real dataset are available (Section VI), not as confirmed findings.

VI. CONCLUSION AND FUTURE SCOPE

A. Summary of the Study

This paper proposed a conceptual architecture for a readiness-aware internship recommendation and skill-gap analysis system, motivated by the observation that most existing internship and job recommenders (Section II) prioritise opportunity matching while giving limited structured attention to whether a candidate is actually ready for a role, or to explaining what stands between them and readiness. The proposed five-layer design (Section III) separates skill extraction, gap analysis, readiness scoring, and explanation into distinct, inspectable modules, and Section IV maps each layer onto a concrete, though not yet built, technology stack.

B. Limitations

The primary limitation of this work is that it has not yet been empirically validated: the evaluation framework in Section V remains a proposal rather than a set of measured results. Two further limitations are worth noting at this conceptual stage.

First, the resume-parsing stage of the Preprocessing and Feature Extraction Layer will depend on text extraction from uploaded documents, and is likely to be less reliable on resumes that rely heavily on non-standard layouts or graphical elements rather than linear text. Second, the learning resources referenced by the resource-mapping component (Section III.E) are intended to be drawn from a curated, periodically updated mapping between skills and resources; without an ongoing maintenance process, this mapping risks becoming stale relative to current course offerings.

C. Future Scope

Future work would proceed in three immediate directions: developing a working prototype of the modules proposed in Section III; assembling or sourcing a suitable dataset of student profiles and internship postings mapped against a shared skill taxonomy (Section III.C); and conducting a pilot evaluation against the metrics proposed in Section V.A.

Beyond this initial validation, longer-term extensions could include integrating live learning-platform APIs (for example, Coursera or Udemy) so that the resource-mapping component updates dynamically rather than from a static, manually curated list; incorporating generative resume feedback, using a language model to give students section-level suggestions for how their resume would be read by an automated applicant-tracking system; and predictive trend forecasting, using sequence models over historical posting data to indicate which skills are likely to be in demand in future placement cycles.

REFERENCES

[1] F. P. Poncio, “Navigating techniques in job recommender systems on internship profile matching: A systematic review,” *Journal of Research in Innovative Teaching & Learning*, vol. 17, no. 2, pp. 352–367, 2024.

[2] D. Celik Ertugrul and S. Bitirim, “Job recommender systems: A systematic literature review, applications, open issues, and challenges,” Journal of Big Data, vol. 12, no. 1, Art. 140, 2025.

[3] Y. Mashayekhi, N. Li, B. Kang, J. Lijffijt, and T. De Bie, “A challenge-based survey of e-recruitment recommendation systems,” ACM Computing Surveys, vol. 56, no. 10, Art. 252, 2024.

[4] F. Tang, R. Zhu, F. Yao, J. Wang, L. Luo, and B. Li, “Explainable person-job recommendations: Challenges, approaches, and comparative analysis,” Frontiers in Artificial Intelligence, vol. 8, Art. 1660548, 2025.

[5] P. Rajapadiyan and Sumithra Devi K., “Smart internship and skill gap analysis,” International Journal of Creative and Open Research in Engineering and Management, 2026.

[6] A. A. J. Permana and G. A. Pradnyana, "Recommendation systems for internship place using artificial intelligence based on competence," IOP Conference Series: Materials Science and Engineering, vol. 1165, 2019.
[7] P. Ajay, K. Arjun, B. Nischay, and R. Lakshmi, "AI-based internship recommendation engine for PM Internship Scheme," International Journal for Innovative Research in Technology, vol. 12, no. 12, pp. 2947–2955, 2026.
[8] "Smart internship recommender," International Research Journal of Modernization in Engineering, Technology and Science (IRJMETS).
[9] "AI-based internship recommendation system," ResearchGate preprint, 2025.
[10] "CareerBERT: Matching resumes to ESCO jobs in a shared embedding space for generic job recommendations," arXiv:2503.02056, 2025.
[11] N. K. Seera and M. Chauhan, "Predicting student employability using machine learning techniques: Performance analysis and comparison," in Revitalizing Student Skills for Workforce Preparation, Hershey, PA: IGI Global Scientific Publishing, 2025, pp. 217–240.
[12] "Next Hire: A data-driven student employability prediction system," International Journal of Engineering Research & Technology (IJERT), vol. 15, no. 4, 2026.
[13] "Design and implementation of an AI-based career counsellor web application," World Journal of Advanced Research and Reviews (WJARR), 2026.
[14] "Career guidance system using decision tree, random forest, and Naive Bayes algorithm," International Journal of Science, Technology and *Society*, Science Publishing Group, 2025.