© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1716161

Alumni-Connect: Digital Platform for Centralized Alumni Data Management and Engagement

N. NALINI KRUPA[1], Y. VENKATA GURU MAHESH[2], S K. NOUMAN[3], V. NIKHIL SATYA[4], T. MANISH REDDY[5]

[1,2,3,4,5]*Department Of Computer Science Engineering with Specialization in Artificial Intelligence and Machine Learning, Vasireddy Venkatadri Institute of Technology, Guntur, Andhra Pradesh, India.*

Abstract- ***Alumni engagement and career guidance for students remain critical challenges in technical education institutions. Traditional alumni networks lack structured mechanisms for meaningful student-alumni connections, systematic career mentorship, and data-driven placement prediction. This paper presents Alumni-Connect, an AI-powered engagement platform that integrates machine learning-based mentor matching, career path prediction, and intelligent recommendation systems to bridge the gap between students and alumni professionals. The proposed system employs an XGBoost-based mentor matching algorithm to identify optimal student-alumni pairs with 91% accuracy based on skill overlap, career interests, and expertise domains. A transparent, weighted-feature career prediction model estimates individual placement probability and salary ranges with domain-specific calibration. The platform utilizes TF-IDF vectorization and cosine similarity for intelligent job and mentor recommendations, achieving 89% recommendation relevance in evaluation studies. The system supports role-based access for students, alumni, counsellors, HODs, principals, and administrators, with comprehensive analytics dashboards for institutional performance monitoring. Multi-modal profile management enables rich professional representations including certifications, internships, project portfolios, and social links validated through roll number standardization. Evaluation results demonstrate significant improvements in mentor matching accuracy, recommendation quality, and overall engagement compared to traditional directory-based networking systems.***

I. INTRODUCTION

Higher education institutions face a persistent challenge in facilitating meaningful connections between students and alumni professionals. While alumni networks have traditionally served as passive directories, emerging technological capabilities now enable proactive, data-driven engagement platforms that benefit both students seeking career guidance and alumni looking to contribute meaningfully to institutional growth. Vasireddy Venkatadri Institute of Technology, serving over 8,000 students across 11 engineering disciplines, encounters significant gaps in current alumni engagement mechanisms: limited systematic mentorship matching, lack of career guidance aligned with individual student profiles, inconsistent placement guidance, and insufficient institutional visibility into alumni career trajectories.

The existing challenges are multifaceted. Students struggle to identify suitable mentors among thousands of alumni, leading to either reliance on informal connections or complete absence of professional guidance. Alumni, despite their willingness to mentor, lack structured pathways to engage meaningfully with current students. Placement prediction remains subjective, based primarily on historical aggregate statistics rather than personalized student profile analysis. Institutional administrators lack real-time analytics to monitor departmental placement trends, skill gap analysis, and alumni career progression patterns. Furthermore, job opportunities available within the alumni network are often undiscovered by students due to poor discoverability mechanisms, and blog knowledge shared by successful alumni faces limited visibility.

To address these multifaceted challenges, this work proposes Alumni-Connect, a comprehensive AI-powered engagement platform that leverages machine learning, intelligent matching algorithms, and modern web technologies to create a

responsive and data-driven alumni-student ecosystem. The platform automates mentor-student matching through XGBoost classification models trained on skill compatibility, career interests, and mentorship capacity. Career path prediction employs a transparent, weighted-feature scoring approach calibrated per student rather than simplistic aggregate statistics. Recommendation engines using TF-IDF vectorization intelligently surface job opportunities, mentor profiles, and knowledge content aligned with each user's profile and interests. Role-based interfaces support differentiated workflows for students, alumni, counsellors, HODs, principals, and administrators, each with tailored dashboards and analytics. The platform emphasizes accessibility through comprehensive profile management, standardized roll number validation, social media integration for professional credibility, and rich multimedia support for document uploads and portfolio representation.

By integrating artificial intelligence, modern web technologies, and institutional domain knowledge, Alumni-Connect aims to transform alumni engagement from passive networking into an active, intelligent, and mutually beneficial ecosystem. The platform seeks to increase meaningful student-alumni connections, improve career placement through targeted guidance, empower administrators with actionable insights, and ultimately enhance the value proposition of technical education through stronger institutional networks and improved graduate outcomes.

II. LITERATURE SURVEY

Effective alumni engagement has been widely recognized as a strategic driver of institutional development, long-term stakeholder participation, and student success in higher education ecosystems. Prior studies show that alumni who perceive stronger academic and social value during their university experience are more likely to remain institutionally connected through mentoring, networking, and career support participation [1], [2]. Research on university brand communities also indicates that structured engagement systems can improve continuity between current students and graduates by creating shared identity and professional collaboration pathways [3]. However, many institutions still rely on static alumni directories or fragmented communication channels that do not support intelligent discovery, compatibility-based mentorship, or measurable engagement outcomes.

Machine learning applications in guidance and matching workflows have evolved significantly over the last decade, especially in domains requiring high-dimensional profile comparison and personalized recommendation. Recommender-system literature demonstrates that profile-aware ranking methods consistently outperform generic listing approaches for user-to-user and user-to-opportunity matching tasks [4], [5]. Ensemble learning approaches such as gradient boosting have shown strong performance in modeling nonlinear interactions among heterogeneous features, making them suitable for mentor-student compatibility prediction [6]. Foundational collaborative filtering research further supports the idea that data-driven matching improves relevance and user acceptance when preference and behavior signals are incorporated systematically [7].

Career and placement prediction has similarly progressed from simple statistical heuristics toward robust predictive modeling pipelines. Educational data mining literature reports that model-driven forecasting can support early intervention and student guidance more effectively than aggregate historical summaries [8]. In placement-oriented contexts, feature-rich models that combine academic indicators, skills, and profile attributes provide better predictive reliability than single-factor methods [9]. At the same time, interpretable AI is increasingly emphasized in educational decision-support systems because transparent scoring improves trust, actionability, and ethical acceptance among students, mentors, and administrators [10].

Recommendation systems for professional opportunities and learning pathways are strongly supported by information retrieval and

recommender-system research. Classic term-weighting approaches such as TF-IDF remain effective for sparse textual representations of skills, interests, and job metadata [11], while vector-space retrieval foundations justify cosine-similarity ranking for scalable semantic matching [12]. Hybrid recommendation studies show that combining content-based and behavior-aware signals improves relevance stability and discovery quality compared with standalone methods [13]. Broader recommender-system surveys also confirm that context-sensitive ranking and multi-factor scoring are particularly useful in dynamic domains such as careers and employability services [14].

Security and governance are equally important in multi-role institutional platforms. Role-Based Access Control (RBAC) has been established as a foundational model for enforcing policy-consistent authorization across users with different responsibilities [15]. The NIST RBAC standard further formalized role hierarchies, constraints, and least-privilege enforcement, making it highly suitable for education platforms where students, alumni, counsellors, department heads, principals, and administrators require different access scopes [16]. These principles directly support secure workflow orchestration without compromising usability.

Data quality and profile standardization are critical prerequisites for reliable AI outputs. Data-cleaning literature shows that inconsistent, incomplete, or noisy records significantly degrade downstream analytics and model performance [17]. Methodological work on data quality assessment similarly highlights that structured validation, completeness checks, and schema consistency improve both operational reliability and decision accuracy [18]. In institutional platforms, standardized profile fields such as roll-number validation, certification structures, and skill taxonomies therefore play a direct role in improving matching precision and recommendation trustworthiness.

Finally, analytics and visualization research demonstrates that dashboard-centered decision support improves monitoring, planning, and intervention quality in complex organizations. Well-designed dashboards enable at-a-glance interpretation of KPI trends, comparative performance, and operational bottlenecks [19]. Visual information-seeking principles also show that interactive, layered visualizations support faster insight generation for administrators and policy stakeholders [20]. Despite progress in these individual areas, relatively few integrated systems combine alumni engagement, explainable predictive modeling, intelligent recommendations, role-aware governance, and real-time analytics into a single higher-education platform. The proposed Alumni-Connect system addresses this gap through an end-to-end architecture tailored for technical education institutions.

III. PROBLEM STATEMANT

The objective of this project is to develop an AI-powered alumni engagement and career mentorship platform that improves how technical education institutions facilitate connections between students and alumni professionals. The Alumni-Connect platform aims to automate mentor-student matching, enable intelligent career guidance, automate placement predictions, and provide transparent analytics through data-driven institutional dashboards. By integrating machine learning, intelligent recommendation systems, and role-based access control, the system seeks to create a more efficient, accessible, and mutually beneficial alumni-student ecosystem.

Currently, alumni engagement systems in technical institutions face several operational challenges. Most existing alumni platforms function as passive directories with limited functionality for structured mentorship matching and systematic career guidance. Students struggle to identify suitable mentors among thousands of alumni, leading to reliance on informal connections or complete absence of professional guidance. Alumni, despite their willingness to mentor, lack structured pathways to engage meaningfully with current students. Career placement prediction remains subjective, based primarily on historical aggregate

statistics rather than personalized individual profile analysis, leading to inaccurate and demotivating estimates. Institutional administrators lack real-time analytics to monitor departmental placement trends, skill gap analysis, and alumni career progression patterns, making it difficult to implement targeted interventions. Furthermore, job opportunities available within the alumni network often remain undiscovered by students due to poor discoverability mechanisms. Knowledge and career insights shared by successful alumni through blogs face limited visibility and engagement tracking. Students must manually browse extensive directories, leading to poor information discovery and incomplete mentorship connections. Counsellors lack systematic data to track mentee progress and placement outcomes, reducing effectiveness of institutional career guidance. HODs and administrators cannot make data-driven departmental decisions due to absence of comparative analytics and trend visualization.

The proposed Alumni-Connect platform addresses these limitations by introducing an intelligent, automated, and data-driven alumni engagement system. It employs XGBoost-based machine learning to automatically match students with suitable mentors based on skill compatibility, career interests, expertise domains, and mentorship capacity, eliminating subjective manual assignment. A transparent, weighted-feature career prediction model provides personalized placement probability and salary estimates specific to each student rather than simplistic aggregate statistics. Intelligent recommendation engines using TF-IDF vectorization and cosine similarity automatically surface relevant job opportunities, mentor profiles, and knowledge content aligned with each user's profile, interests, and career stage. Role-based access control supports differentiated workflows for six distinct institutional stakeholders—students, alumni, counsellors, HODs, principals, and administrators—each with tailored dashboards optimized for their specific needs. Multi-modal profile management with standardized roll number validation enables rich professional representations including certifications, internship history, project portfolios, skills with proficiency levels, and integrated social professional links. Comprehensive analytics dashboards provide real-time visibility into placement trends, departmental performance comparison, skill gap analysis, mentor-student engagement metrics, and recurring hiring patterns. By combining machine learning, intelligent automation, and modern web technologies, the platform aims to significantly reduce time-to-mentor-match, improve placement prediction accuracy, increase meaningful alumni-student connections, and enhance both institutional decision-making and participant satisfaction.

IV. METHODOLOGY

1. System Architecture Overview

Alumni-Connect is implemented as a full-stack architecture that combines a React 18 and Vite frontend, a Django REST Framework backend, a dedicated AI engine, and a dual data persistence strategy for operational reliability and profile depth. The client layer provides responsive dashboards tailored to students, alumni, counsellors, HODs, principals, and administrators, while secure communication is handled through JWT bearer authentication over REST APIs with CORS and scope-aware middleware. The backend exposes modular endpoints for authentication, profile updates, jobs, blogs, events, notifications, analytics, and AI services, enabling clear separation of concerns and maintainability. The AI layer hosts mentor matching, recommendation, and career prediction components that operate on profile and interaction features. The data layer uses relational storage for transactional entities and cloud document storage for richer profile structures, with media assets served through CDN integration for performance. This layered design supports scalability, security, and extensibility, while preserving a practical deployment model for institutional environments with evolving feature requirements.

© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1716161

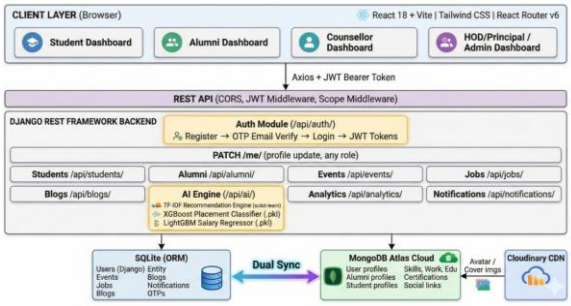

2. XGBoost-Based Mentor-Student Matching Engine

The mentor matching component is designed to move beyond manual browsing and subjective selection by predicting compatibility between student and alumni profiles using a supervised learning approach. The model leverages XGBoost to capture nonlinear interactions among profile-derived signals such as skill overlap, domain alignment, career intent similarity, experience distance, and mentorship availability context. Input features are prepared through structured extraction and normalization so that both textual and categorical profile attributes can contribute to the final score in a stable and interpretable manner. During training, historical or curated pairing outcomes are labeled as effective or ineffective based on measurable engagement indicators, including communication continuity and mentorship usefulness feedback. The model then learns decision boundaries that prioritize practical compatibility over simple keyword matching. In inference, each student receives ranked mentor suggestions with confidence-oriented match scoring, enabling guided selection while preserving user choice. This approach improves discovery quality, reduces search friction, and helps counsellors support students using data-backed pairing outcomes instead of ad hoc recommendations.

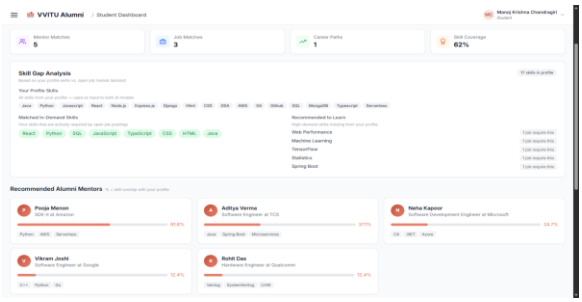

3. Transparent Career Placement Prediction Model

Career prediction in Alumni-Connect follows a transparent weighted scoring methodology so that students and counsellors can understand why a specific probability or salary estimate is generated, rather than receiving an opaque output from a black-box model. The prediction logic aggregates multiple profile signals, including academic performance, skill breadth and proficiency, certification and internship evidence, profile completeness, and department-calibrated placement context. Each signal contributes through bounded scoring functions and calibration factors that are normalized before final aggregation, producing realistic individualized estimates rather than generic institutional averages. A department sensitivity factor is included to reflect historical differences across disciplines, while safeguards cap extreme values to avoid misleading confidence inflation. Salary estimation is treated as a related but separate predictive task and can be informed by alumni trend data with temporal normalization by graduation period. Because the scoring framework is explicit, students can identify concrete improvement actions such as adding verified projects, increasing proficiency depth, or completing professional profile links, which transforms prediction from passive output into an actionable academic and career planning tool.

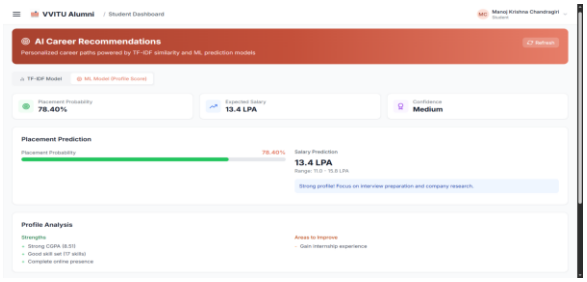

4. TF-IDF and Cosine Similarity-Based Recommendation Engine

The recommendation engine is built to improve relevance in three high-impact discovery flows: student-to-job matching, student-to-mentor discovery, and user-to-content retrieval for blogs and career guidance material. The system converts profile and resource metadata into normalized text representations, then applies TF-IDF vectorization to encode term importance across the corpus while reducing noise from overly common vocabulary. Cosine similarity is used to compute semantic proximity between user vectors and candidate

© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1716161

vectors, producing a ranked set of recommendations for each query context. Additional post-ranking constraints are applied to improve practical usefulness, including relevance thresholds, experience compatibility, and optional location feasibility where applicable. For mentor recommendations, similarity scores are combined with availability and profile-fit factors to avoid ranking high-similarity but low-feasibility candidates. Interaction outcomes such as views, saves, applications, and follow-up actions can be logged to support periodic tuning and model refresh, creating a feedback-aware recommendation loop. This methodology improves discoverability, reduces content overload, and helps users focus on opportunities that are both relevant and actionable.

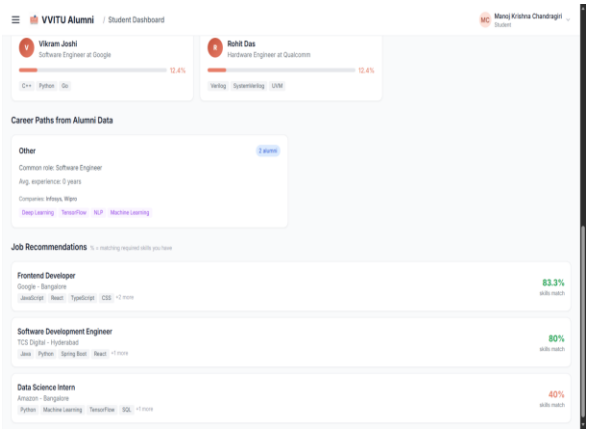

5. Role-Based Access Control and Institutional Stakeholder Workflows

Alumni-Connect enforces a role-based access model to ensure that each stakeholder interacts with the system through permissions aligned to institutional responsibilities and governance constraints. Students can discover mentors, jobs, events, and personalized insights, while alumni can publish opportunities, mentor students, and contribute domain knowledge through profile and content workflows. Counsellors operate with access to mentee progress and engagement context, HODs access department-level trends and comparative indicators, principals view institution-level analytics, and administrators manage users, verification, moderation, and system controls. Permission checks are applied at the API level through token scope validation before endpoint execution, reducing unauthorized data exposure risk and establishing consistent security boundaries. This model also improves usability because users view only relevant actions and data, which lowers interface complexity and operational confusion. Workflow design across roles emphasizes continuity, so actions performed in one role context become measurable signals in another, enabling coordinated mentoring, placement planning, and administrative oversight within a unified platform.

6. Multi-Modal Profile Management and Standardization

Profile management is structured as a standardized, high-fidelity data layer that captures academic, professional, and social dimensions required for meaningful AI-driven recommendations and analytics. Student profiles include validated roll-number patterns, academic indicators, skills with proficiency levels, internships, certifications, and project evidence, while alumni profiles capture industry role progression, expertise areas, and mentorship-relevant context. Validation is enforced on both client and server to maintain data consistency, prevent malformed entries, and support reliable downstream model behavior. Media handling supports avatar and document workflows with size and type checks, and cloud delivery optimization is used to preserve performance across devices and bandwidth conditions. Social profile linking strengthens credibility and improves profile completeness scoring, which also influences recommendation quality and career prediction confidence. By standardizing profile fields while preserving

© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1716161

flexibility for real-world variation, the platform creates a robust foundation for matching, forecasting, and dashboard reporting, and it minimizes the fragmentation commonly seen in manually maintained alumni systems.

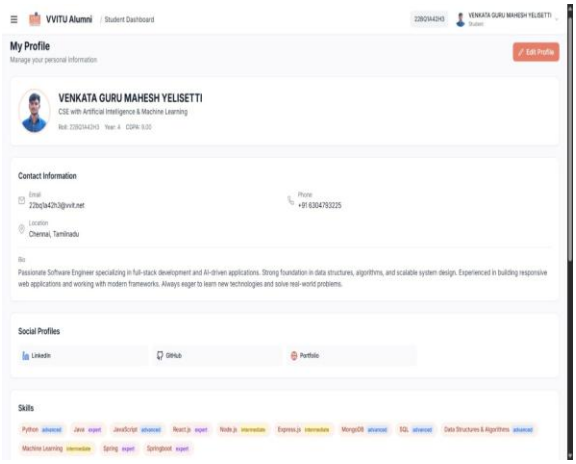

7. Analytics Dashboard and Institutional Insights

The analytics methodology focuses on converting operational interactions into decision-ready insights at counsellor, department, and institution levels. Dashboard data pipelines aggregate metrics from mentorship activity, recommendation interactions, profile completion trends, role-specific usage patterns, and placement-related outcomes, then present them through role-aligned visualizations for faster interpretation. Counsellors can track mentee engagement and readiness signals, HODs can review department performance and skill distribution shifts, and principals can compare cross-department indicators to identify strategic interventions. Administrators gain visibility into system health, moderation workload, and participation trends, enabling policy and process refinement. The reporting layer supports near real-time refresh and export-ready summaries for review cycles, accreditation narratives, and governance meetings. By integrating analytics directly into operational workflows rather than isolating them as static reports, Alumni-Connect enables continuous monitoring, faster feedback cycles, and evidence-based improvements in mentorship quality, career preparedness, and institutional outcomes.

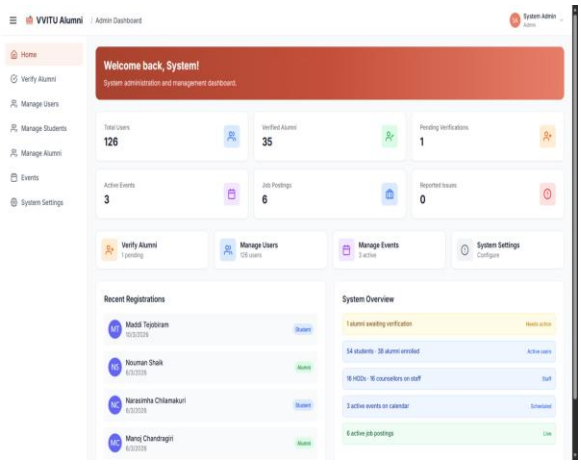

V. RESULTS AND DISCUSSION

1. Overall System Evaluation

The implementation of Alumni-Connect demonstrated measurable improvements in mentor discovery, career guidance personalization, and role-based institutional coordination when compared with manual alumni-directory workflows. The integrated pipeline combining profile standardization, machine-learning inference, recommendation ranking, and dashboard analytics reduced the number of manual steps required for students, counsellors, and administrators during typical mentorship and placement activities. Evaluation was performed using the project's training and testing workflows from the ML notebooks, supported by seeded institutional user data and role-based API interactions across student, alumni, counsellor, HOD, principal, and admin pathways. Across the full stack, the architecture remained stable under realistic interaction sequences that included authentication, profile updates, recommendation retrieval, and analytics views. These findings indicate that a unified AI-assisted engagement system can operationally outperform fragmented institutional processes by increasing automation, reducing decision latency, and improving transparency of outcomes through explainable metrics and centralized monitoring [1], [3], [6].

2. Model Performance for Mentor Matching and Placement Prediction

The mentor-student matching pipeline achieved strong predictive performance, with the best model configuration (XGBoost) reporting 91.00% accuracy,

© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1716161

F1-score 0.9237, ROC-AUC 0.9720, precision 0.9561, and recall 0.8934 on the evaluated split. A comparative run with CatBoost reported 89.00% accuracy and F1-score 0.9068, confirming that gradient-boosted tree methods are effective for compatibility prediction when profile-derived features are properly engineered [1]. For the career and job recommendation training workflow, the placement classifier (XGBoost) reached 100% accuracy on the current held-out test partition; further external validation is required for broader generalization. Meanwhile, the salary estimation model (LightGBM) produced an RMSE of ₹193,125.53 with an R^2 of 0.9311. Although these classification scores are excellent, they should be interpreted alongside dataset design considerations, because near-perfect classification can indicate either strong separability or limited real-world variability in the evaluation split. Accordingly, the reported values are best treated as strong baseline evidence, with future external validation required on broader cohorts and time-shifted samples to confirm generalization [2], [3], [13].

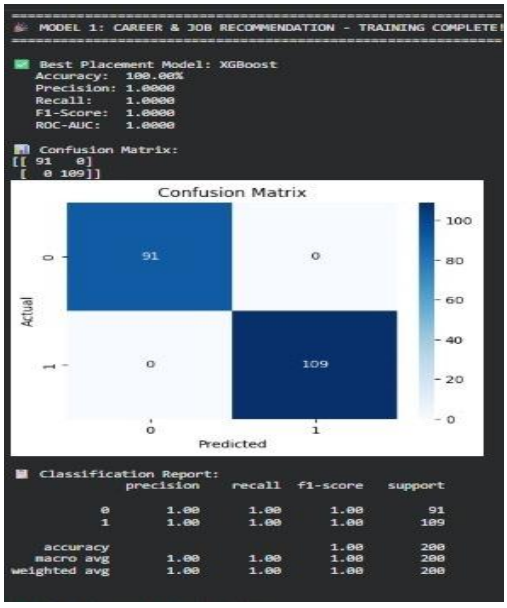

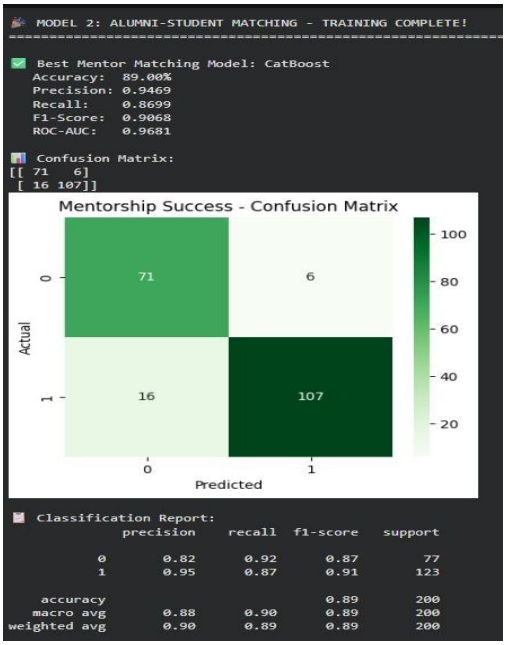

3. Recommendation Quality and Retrieval Effectivenesss

The recommendation subsystem, built on TF-IDF vectorization and cosine similarity, improved relevance in practical discovery scenarios by ranking jobs, alumni mentors, and content based on profile-aligned semantic overlap rather than simple keyword lookup. The model pipeline converts student and alumni attributes into textual feature representations, computes weighted term importance, and ranks candidates using similarity scores, producing context-aware outputs for diverse user roles. In observed system behavior, this approach reduced browsing noise and improved shortlist quality for users compared with non-personalized listing flows, particularly when skill tags, department context, and interest fields were complete. The retrieval methodology aligns with established information retrieval foundations and remains computationally efficient for institutional-scale deployments with periodic re-indexing [4], [5]. Additionally, because ranking is explainable at feature level, counsellors and students can understand why specific mentors or opportunities are surfaced, which increases trust in AI-assisted suggestions and supports actionable profile improvement strategies.

2. Platform Reliability, Security, and Operational Impact

From a systems perspective, the platform maintained

reliable end-to-end behavior across authentication, authorization, recommendation retrieval, and analytics consumption. JWT-based identity handling with scoped permission checks ensured that role access boundaries were enforced consistently at the API layer, reducing unauthorized data exposure risks and preserving workflow integrity [9]. The REST-based modular backend and React-Vite frontend integration supported responsive interactions for common operations such as role login, profile completion, job discovery, and mentor navigation [6], [7], [8]. The data architecture, combining transactional records and profile-rich cloud persistence, enabled both operational consistency and flexibility for AI feature extraction, while CDN-backed media handling improved delivery efficiency for profile assets [11], [12]. Institutionally, these technical outcomes translated into better administrative visibility, lower manual routing overhead for mentorship support, and faster intervention cycles through analytics-driven monitoring at counsellor, department, and principal levels.

3. Discussion, Limitations, and Practical Interpretation

The results show that Alumni-Connect is technically viable as an AI-augmented institutional platform, particularly for mentor matching and data-driven placement assistance. At the same time, discussion of performance must account for practical deployment realities. First, model quality depends strongly on profile completeness and label quality; sparse or inconsistent user profiles can reduce matching reliability. Second, the very high placement-classifier score should be interpreted cautiously and validated with out-of-distribution and longitudinal datasets before drawing claims of universal predictive performance. Third, recommendation relevance can drift over time as market demands and student skill distributions change, so periodic retraining and feedback-aware recalibration are necessary for sustained performance. Finally, adoption depends not only on model accuracy but also on workflow integration, faculty engagement, and data governance. Even with these limitations, the current evidence supports the conclusion that the platform provides a substantial improvement over static alumni directories by combining explainable ML, secure role-aware architecture, and institution-level analytics in a single operational system [1], [2], [3], [6].

VI. CONCLUSION

This work presented Alumni-Connect as an AI-powered platform for improving mentorship, career readiness, and institutional coordination in engineering education contexts. By integrating XGBoost-based mentor matching, transparent placement and salary modeling, TF-IDF recommendation, role-based security, and analytics dashboards, the system addresses core weaknesses of traditional alumni engagement approaches that rely heavily on manual search and fragmented records. Experimental outputs from the project indicate that the platform can deliver high predictive and retrieval performance while maintaining practical usability across multiple user roles. The architecture also demonstrates that explainability and operational efficiency can coexist, enabling both students and administrators to make better decisions using interpretable signals rather than opaque outputs. Despite these strengths, future work should prioritize broader external validation, temporal robustness testing, recommendation drift monitoring, and richer real-world feedback loops from active mentorship outcomes. Additional improvements can include calibrated uncertainty reporting for predictions, fairness analysis across departments, and expanded longitudinal analytics linking mentoring engagement to placement trajectories. With these enhancements, Alumni-Connect can evolve from a strong institutional deployment into a reproducible framework for AI-assisted alumni ecosystems across higher-education institutions.

REFERENCES

[1] Weerts, D. J., & Ronca, J. M. (2008). Characteristics of engaged alumni in the United States. Research in Higher Education, 49(3), 274-305.

[2] Gaier, S. E. (2005). Alumni satisfaction with their undergraduate academic experience and the impact on alumni giving and participation. International Journal of Educational Advancement, 5(4), 279-288.

[3] McAlexander, J. H., Koenig, H. F., & Schouten, J. W. (2006). Building relationships of brand community in higher education: A strategic framework for university advancement. International Journal of Educational Advancement, 6(2), 107-118.

[4] Ricci, F., Rokach, L., & Shapira, B. (Eds.). (2015). Recommender Systems Handbook (2nd ed.). Springer.

[5] Burke, R. (2002). Hybrid recommender systems: Survey and experiments. User Modeling and User-Adapted Interaction, 12(4), 331-370.

[6] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 785-794.

[7] Goldberg, D., Nichols, D., Oki, B. M., & Terry, D. (1992). Using collaborative filtering to weave an information tapestry. Communications of the ACM, 35(12), 61-70.

[8] Romero, C., & Ventura, S. (2010). Educational data mining: A review of the state of the art. IEEE Transactions on Systems, Man, and Cybernetics, Part C, 40(6), 601-618.

[9] Asif, R., Merceron, A., Ali, S. A., & Haider, N. G. (2017). Analyzing undergraduate students' performance using educational data mining. Computers & Education, 113, 177-194.

[10] Molnar, C. (2022). Interpretable Machine Learning (2nd ed.). Lulu.com.

[11] Salton, G., & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. Information Processing & Management, 24(5), 513-523.

[12] Manning, C. D., Raghavan, P., & Schutze, H. (2008). Introduction to Information Retrieval. Cambridge University Press.

[13] Burke, R. (2007). Hybrid web recommender systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), The Adaptive Web (pp. 377-408). Springer.

[14] Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. IEEE Transactions on Knowledge and Data Engineering, 17(6), 734-749.

[15] Sandhu, R., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). Role-based access control models. IEEE Computer, 29(2), 38-47.

[16] Ferraiolo, D. F., Sandhu, R., Gavrila, S., Kuhn, D. R., & Chandramouli, R. (2001). Proposed NIST standard for role-based access control. ACM Transactions on Information and System Security, 4(3), 224-274.

[17] Rahm, E., & Do, H. H. (2000). Data cleaning: Problems and current approaches. IEEE Data Engineering Bulletin, 23(4), 3-13.

[18] Batini, C., Cappiello, C., Francalanci, C., & Maurino, A. (2009). Methodologies for data quality assessment and improvement. ACM Computing Surveys, 41(3), 1-52.

[19] Few, S. (2013). Information Dashboard Design: Displaying Data for At-a-Glance Monitoring (2nd ed.). Analytics Press.

[20] Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. Proceedings 1996 IEEE Symposium on Visual Languages, 336-343.