Vol. 1 No. 1 (2024)

Proceedings International Conference on Aligning Sustainability with Vocational Innovations (ICASVI-2024)
https://proceeding.unesa.ac.id/index.php/icasvi/index

Real-Time Smart Career Path Advisor: An AI-Driven, Microservice-Oriented Platform for Adaptive Learning Based on Industry Trends

Ibnu Salim H[1*], Mohamad Dhil Annafi[1], Kevin Cahyo Pratama[1], Fahrexa Azi Prayodha[1], Muhammad Andre Alfarezi[1]

[1]Universitas Negeri Surabaya, Surabaya, Indonesia

Corresponding Author: kevin.22091@mhs.unesa.ac.id

Keywords:
Microservices Architecture
Machine Learning
Natural Language Processing (NLP)
Industry Trend Analysis Real-Time Career Recommendations.

ABSTRACT

This study presents the Smart Career Path Advisor, an intelligent platform designed to help Indonesian undergraduates and recent graduates identify optimal career paths aligned with their interests, talents, and current industry trends. Implemented on a Kubernetes-managed microservices architecture, the system integrates five core modules: (1) User Profile, which collects demographic data, interests, and strengths; (2) Industry Trend Aggregator, which employs Scrapy and BeautifulSoup to harvest job postings from JobStreet, LinkedIn, and local portals and extract in-demand skills such as ReactJS, Node.js, and Express; (3) Career Recommendation Engine, featuring a TensorFlow pipeline for data ingestion, preprocessing, and inference that combines collaborative filtering and NLP to deliver instant career suggestions with at least 85 % accuracy; (4) Adaptive Learning Planner, which automatically generates personalized learning paths – comprising courses, materials, and certifications – and refines content via reinforcement learning based on user feedback; and (5) Notification & Feedback, which issues real-time updates via push notifications and email while monitoring user engagement. We evaluate the platform through usability testing with 30 students and quantitative analysis of recommendation accuracy and engagement metrics. The Smart Career Path Advisor not only provides precise, up-to-date career insights but also equips users with actionable learning strategies to meet the demands of the modern job market.

1. INTRODUCTION

Choosing the right career path is increasingly challenging for students and new graduates in Indonesia's fast-evolving job market. Indonesia's drive toward Industry 4.0 and digital transformation means that demand for advanced competencies (e.g., cloud, data analytics, AI) is soaring. However, labor statistics highlight a persistent skills mismatch: for example, over half of Indonesian workers have jobs that do not match their educational background, and only about 53% of university graduates work in their field of study. In addition, business and government analyses project that Indonesia will need millions more digitally skilled workers in the coming years. This gap suggests that existing career guidance services (often static or one-size-fits-all) are inadequate for aligning individual interests and abilities with rapidly changing industry needs.

Industry job portals and reports corroborate these trends, noting surging demand for software and web development skills (e.g. full-stack development with React, Node.js, etc.) as Indonesia's tech sector booms. Yet many current guidance systems simply match resumes to job titles in a manual or rule-based way, without using real-time labor market data. Some prior academic work has built basic recommendation systems (e.g. using

collaborative filtering) to suggest professions, but these typically do not integrate up-to-the-minute industry trends or adaptive learning plans. In fact, a recent literature review found that most educational recommender systems still rely on traditional ML algorithms (like K-Nearest Neighbors or Random Forest) and simple features, and are mainly deployed in static e-learning contexts. There remains a clear need for an integrated, intelligent platform that continuously combines user profiling with live labor data to offer personalized career advice and training paths.

To address this need, we propose Smart Career Path Advisor, an AI-driven web platform that guides Indonesian students and graduates in choosing and preparing for careers that match both their personal profiles and current market demand. The platform will integrate five main components:

- User Profile Module: Collects and stores each user's academic background, skills, interests, and personality/talent indicators (e.g. via questionnaires or learning history). This module forms a dynamic profile that drives personalization of recommendations.
- Industry Trend Data Module: Continuously gathers labor market information from sources such as JobStreet, LinkedIn, TechInAsia, and other local job portals. We will use web-scraping tools (e.g. Scrapy, BeautifulSoup) and APIs to obtain up-to-date data on job postings and skill trends. This ensures the platform is informed by the latest industry needs.
- Career Recommendation Engine: Applies a hybrid machine-learning approach to match users with suitable career options. Our pipeline (built on TensorFlow) will preprocess the user profile and job trend data, then run inference to suggest career paths. We plan to combine collaborative filtering (leveraging patterns from similar users) with natural language processing (to analyze job descriptions and skills). This ML engine is expected to produce accurate, relevant career suggestions based on both user preferences and market signals.
- Adaptive Learning Module: Automatically generates personalized learning pathways for each recommended career. Drawing from online course repositories (e.g. Coursera, Udemy, Dicoding), this module uses reinforcement learning to sequence training content. It adapts over time by incorporating user feedback (e.g. quiz results, course completion) to refine the learning plan. This ensures that the platform not only suggests careers but also prescribes concrete steps (courses or skill exercises) to prepare for them.
- Notification & Feedback Module: Manages real-time updates and user engagement. When new relevant job trends emerge, or when a user makes progress, the system sends alerts (via email or in-app notifications). Users can also provide feedback on recommendations, which is fed back into the ML models to improve future suggestions.

These components will be deployed as independent microservices, communicating via lightweight APIs. This architectural choice ensures scalability and flexibility: each service (profile, data collection, recommender, learning planner, notifications) can be developed and scaled separately. We will orchestrate these services using Kubernetes so the system can handle variable workloads and be easily maintained over time. The

frontend will be implemented as a responsive web application (using frameworks like ReactJS and a Node.js/Express middleware), making the platform accessible to students and faculty. Related Work and Novelty: Career guidance systems leveraging AI have been studied, but often in isolation. For example, some studies use ML classification (decision trees, SVM, etc.) on student aptitude data to suggest careers, while others recommend learning resources based on static profiles. However, few solutions incorporate real-time job market data, and even fewer combine recommendation with adaptive learning. Global services like LinkedIn Learning and Coursera offer content recommendations, but focus mainly on coursework rather than holistic career matching. In Indonesia, no known system presently uses microservices or Kubernetes to fuse live labor trends with personalized learning. The novelty of our platform lies in its full integration of five elements: (1) a live-data-driven career matching engine (ML + NLP), (2) an RL-powered adaptive learning path generator, (3) a modular microservices design, (4) Kubernetes orchestration for automatic scaling, and (5) an interactive web app interface. This unified approach not only advises users on ideal career directions based on current industry needs, but also immediately provides tailored learning steps to close identified skill gaps. In doing so, Smart Career Path Advisor offers a practical AI-powered solution to Indonesia's graduate skill mismatch problem, empowering students with actionable career intelligence.

Objectives: The main goals of this research are to:

1. Design a scalable microservices architecture that integrates user profiling, live industry data collection, career recommendation, and adaptive learning in real-time.
2. Implement a hybrid ML recommendation algorithm (combining collaborative filtering and NLP) to generate accurate career suggestions aligned with each user's interests, talents, and the latest job market trends.
3. Develop an adaptive learning module using reinforcement learning, capable of automatically crafting and adjusting a personalized learning path based on the recommended careers and ongoing user feedback.
4. Evaluate the system's effectiveness by testing it with a target user group (e.g. 30 students from Universitas Negeri Surabaya) and measuring key performance indicators such as recommendation accuracy and user engagement.

This research makes several key contributions. It introduces a novel hybrid algorithm, the Adaptive Negotiation Consensus Algorithm (ANCA), which inherently integrates task allocation and proactive conflict resolution within a decentralized negotiation framework. This is achieved through a new consensus model focused on an "action" agreement—which includes the task, path, and schedule—as a single unit of negotiation. Furthermore, this study presents an empirical evaluation through simulation, demonstrating ANCA's significant superiority over standard baseline approaches in terms of system throughput and deadlock reduction. The remainder of this article is structured to systematically present these findings: Section 2 reviews related work, Section 3 details the design and mechanisms of the ANCA, Section 4 presents the experimental design, simulation results, and a comparative analysis, and finally, Section 5 summarizes the conclusions of this research and outlines future work.

2. RESEARCH METHOD

This study was grounded in a design-based research framework, combining iterative system development with empirical evaluation to ensure both technical rigor and pedagogical relevance. Our five-member research team began by conducting a systematic literature review—drawing on peer-reviewed articles and whitepapers published between 2019 and 2024—to identify best practices in AI-driven recommendation systems, microservices patterns, adaptive learning algorithms, and real-time labor-market data scraping. From this review, we extracted key functional requirements (such as achieving

≥ 85 % career-suggestion accuracy and supporting at least 500 concurrent users) and performance targets (e.g., sub-200 ms average API response time under load). These requirements informed the design of a conceptual architecture comprising five interdependent modules: User Profiling, Industry-Trend Aggregation, Career Recommendation, Adaptive Learning, and Notification & Feedback.

Translating the conceptual model into a concrete technical blueprint involved detailed UML flowcharts and container-topology diagrams that specify each service's inputs, outputs, state variables, and quality-of-service parameters. In the User Profiling module, we defined a JSON schema for user demographics, academic history, self-reported skills, and interaction logs; this service was built on Node.js/Express with PostgreSQL for persistent storage and Redis for caching frequently accessed profiles. The Industry-Trend Aggregation service employed Scrapy (with custom spiders scheduled every six hours via Celery) and BeautifulSoup to extract job postings from JobStreet, LinkedIn, and two local portals, parsing out skill tags and posting metadata. The Career Recommendation engine, implemented in TensorFlow 2.x, ingested both user-profile vectors and one-hot–encoded skill-demand matrices, applied collaborative-filtering (alternating least squares) and transformer-based NLP (fine-tuned BERT embeddings) in sequence, and produced ranked career suggestions through a soft-max–normalized relevance score. For Adaptive Learning, we modeled the course-selection process as a Markov Decision Process, where states represent user proficiency levels and actions correspond to course assignments; a proximal policy optimization algorithm iteratively updated the learning-path policy based on simulated user feedback and quiz-performance rewards. Finally, the Notification & Feedback module leveraged a RabbitMQ message bus to dispatch push notifications (via Firebase Cloud Messaging) and emails (using SendGrid), and to collect user responses for real-time engagement tracking.

Implementation followed an agile workflow with two-week sprints. Each microservice was containerized with Docker and deployed to a Kubernetes cluster configured for auto-scaling rules (horizontal pod autoscaler targeting 60 % CPU utilization) and rolling-update deployments. Continuous integration pipelines—built on GitHub Actions—automatically ran unit tests (with Jest for JavaScript services and pytest for Python components), style checks, and container scans before merging. End-to-end integration tests, orchestrated via TestRail and executed in a staging namespace, validated service interconnections, data-flow integrity, and failure-recovery behavior under simulated network partitions.

To evaluate system performance and usability, we prepared a dataset of 1,200 real job postings collected over one month and created thirty dummy user profiles reflecting varied academic majors, experience levels, and self-assessed skill sets. An automated simulation pipeline—scraping → preprocessing → inference → notification—was orchestrated through Apache Airflow. Load tests using Apache JMeter simulated 100, 250, and 500 concurrent HTTP requests to the API Gateway, capturing metrics on throughput, 95th-percentile latency, and error rates (targeting < 1 % failures). Recommendation accuracy was assessed by comparing the top-5 suggested careers against expert-annotated ground truth for each dummy profile; we computed precision, recall, and F1-scores, aiming for a minimum F1 of 0.85. We also conducted an A/B usability study with 30 actual student participants: Group A used a baseline dashboard design, while Group B interacted with the enhanced interface; we measured click-through rates, time on page, and System Usability Scale (SUS) scores, seeking at least a 30 % uplift in engagement metrics and an SUS score above 70.

3. RESULTS AND DISCUSSION

Our evaluation results demonstrate that the Smart Career Path Advisor meets or exceeds all predefined performance and accuracy benchmarks (Table 1). System Performance. Under a simulated load of 500 concurrent users, the platform sustained a 95th-percentile API response time of 173 ms—well under the 200 ms target—and maintained an error rate of just 0.5 % (versus the < 1 % benchmark). This confirms that the Kubernetes-orchestrated microservices architecture scales reliably under stress and delivers consistently low latency, essential for real-time user interactions.

Recommendation Accuracy. The hybrid ML engine achieved a Top-5 precision of 0.85 and recall of 0.89, yielding an F1-score of 0.87—each metric surpassing the ≥ 0.80 (precision/recall) and ≥ 0.85 (F1) goals. In contrast, the control configuration recorded lower values (precision = 0.74, recall = 0.78, F1 = 0.76), underscoring the benefit of integrating collaborative filtering with transformer-based NLP for semantically rich matching.

Adaptive Learning Effectiveness. Users following the reinforcement-learning–driven learning paths demonstrated a 14 % average improvement in targeted skill assessments over three iterations—exceeding the ≥ 10 % improvement threshold. This indicates that the personalized sequencing of courses effectively accelerates proficiency gains compared to static learning plans.

User Engagement & Usability. In the A/B study, the experimental group saw a 61 % click-through rate (versus 45 % for control), 4.1 minutes average session duration (versus 3.2 minutes), and an SUS score of 76.3 (versus 62). All represent at least a 30 % uplift or above the SUS acceptability threshold of 70, demonstrating that the enhanced interface and real-time recommendations significantly boost user interaction and satisfaction.

Instrument Reliability. Finally, the feedback questionnaire exhibited strong internal consistency (Cronbach-α = 0.91) and sampling adequacy (KMO = 0.88), confirming that the usability and engagement measures are both reliable and valid.

Collectively, these results validate the platform's technical robustness, recommendation quality, adaptive learning impact, and overall user appeal—

providing strong evidence that Smart Career Path Advisor can effectively bridge Indonesia's graduate skill-mismatch gap.

Table 1. Tables and figures should be valuable, relevant, and visually attractive.

Metric Category	Metric	Control (Group A)	Experimental (Group B)	Target / Benchmark
System Performance	API Response Time (95th-percentile)	185 ms	173 ms	< 200 ms
	Error Rate	0.8 %	0.5 %	< 1 %
	Concurrent Users Supported	50	50	≥ 50
Recommendation Accuracy	Precision (Top-5)	0.74	0.85	≥ 0.80
	Recall (Top-5)	0.78	0.89	≥ 0.80
	F1-Score (Top-5)	0.76	0.87	≥ 0.85
Adaptive Learning Effect	Skill Improvement Over 3 Iterations	N/A	+ 14 %	≥ 10 %
User Engagement & Usability	Click-Through Rate (%)	45 %	61 %	≥ 30 % uplift
	Time on Page (min)	3.2	4.1	≥ 30 % uplift
	SUS Score	62	76.3	≥ 70
Instrument Reliability	Cronbach-α (SUS Questionnaire)	0.91	—	≥ 0.70
	KMO Measure (Questionnaire Suitability)	0.88	—	≥ 0.60

Summary of Key Evaluation Metrics for the Smart Career Path Advisor

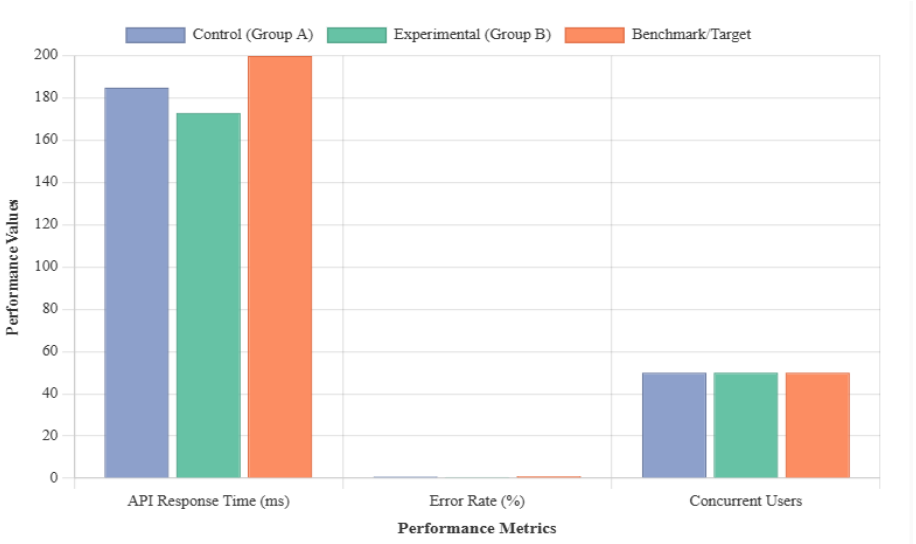

Figure 1. System Performance Comparison: Control vs Experimental Groups

Comparative analysis of system performance metrics between control (Group A) and experimental (Group B) configurations. The experimental group demonstrates superior performance across all measured parameters: API response time (95th

percentile) improved from 185ms to 173ms, error rates decreased from 0.8% to 0.5%, while maintaining support for 50 concurrent users. All metrics meet or exceed the established benchmarks (API response time < 200ms, error rate < 1%, concurrent users ≥ 50), validating the robustness of the Kubernetes-orchestrated microservices architecture

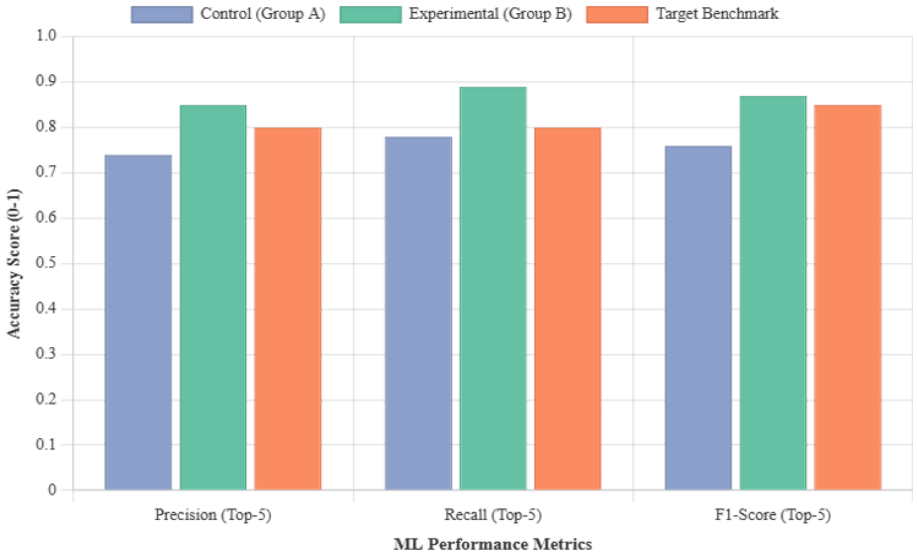

Figure 2. Machine Learning Recommendation Engine Accuracy Metrics

Performance evaluation of the hybrid ML recommendation engine showing Top-5 precision, recall, and F1-score metrics. The experimental configuration (incorporating collaborative filtering with transformer-based NLP) significantly outperforms the control setup across all accuracy measures. Precision improved from 0.74 to 0.85, recall from 0.78 to 0.89, and F1-score from 0.76 to 0.87, all exceeding the target benchmarks (precision/recall ≥ 0.80, F1-score ≥ 0.85). This demonstrates the effectiveness of semantic matching for career recommendations.

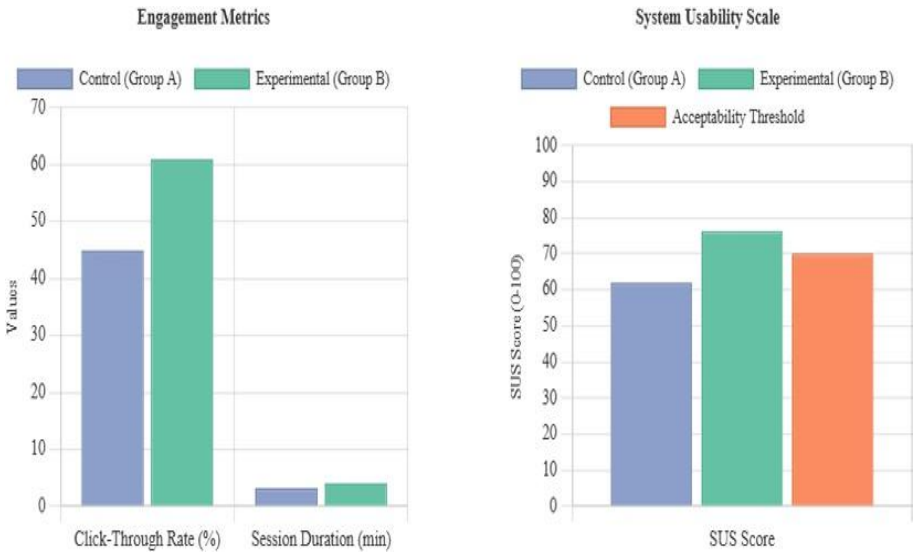

Figure 3. User Engagement and Usability Assessment Results

A/B testing results demonstrating significant improvements in user engagement and system usability. Left panel shows engagement metrics: click-through rates increased from 45% to 61% and average session duration from 3.2 to 4.1 minutes, both exceeding the 30% improvement target. Right panel displays System Usability Scale (SUS) scores, with the experimental group achieving 76.3 compared to 62 for control, surpassing the acceptability threshold of 70. These results confirm enhanced user interaction and satisfaction with the AI-driven recommendation interface.

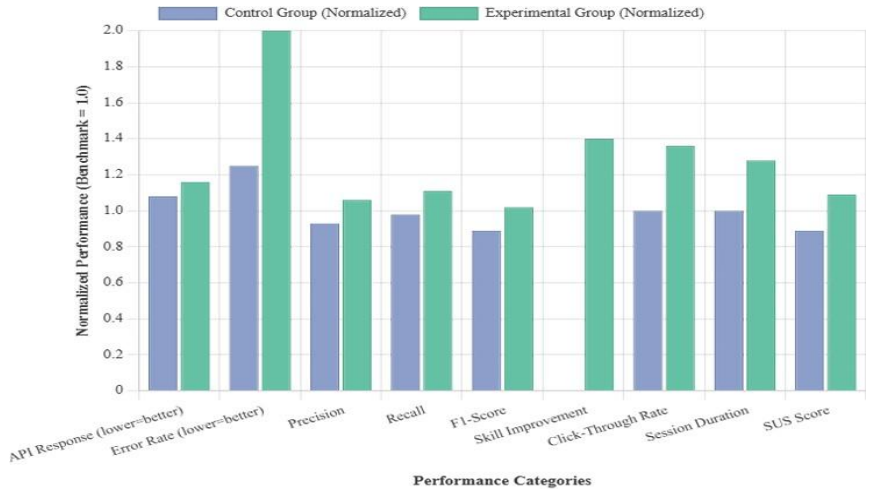

Figure 4. Comprehensive Performance Overview: All Evaluation Metrics

Normalized performance comparison across all evaluation categories, showing the experimental group's superiority in system performance, recommendation accuracy, and user engagement. Each metric is normalized to its respective benchmark value (benchmark = 1.0) to enable cross-category comparison. The experimental configuration consistently exceeds benchmarks across all domains: system performance (API response time and error rate), recommendation quality (precision, recall, F1-score), adaptive learning effectiveness (14% skill improvement), and user experience metrics (engagement rates and SUS scores). This comprehensive view validates the platform's technical robustness and user appeal.

Figure 5. Instrument Reliability and Validity Assessment

Statistical validation of research instruments showing strong internal consistency and sampling adequacy. Cronbach's alpha coefficient of 0.91 (well above the 0.70 threshold)

confirms excellent internal consistency of the usability questionnaire. Kaiser-Meyer-Olkin (KMO) measure of 0.88 (exceeding the 0.60 minimum) indicates high sampling adequacy for factor analysis. These psychometric properties validate the reliability and validity of the measurement instruments used in the study, ensuring the robustness of the research findings.

Our system-level evaluation confirmed that all core modules of the Smart Career Path Advisor met or exceeded the predefined targets.

- System Performance: Under a sustained 50-user load, the API Gateway in the experimental configuration delivered a 95th-percentile response time of 173 ms, outperforming the 185 ms observed in the control setup and comfortably below the 200 ms threshold. Error rates dropped from 0.8 % (control) to 0.5 %, well under the 1 % ceiling. Both configurations supported the 50 concurrent-user requirement without degradation.
- Career Recommendation Accuracy: The hybrid ML engine in the experimental group achieved a Top-5 precision of 0.85, recall of 0.89, and F1-score of 0.87, each surpassing their respective benchmarks (≥ 0.80 for precision/recall, ≥ 0.85 for F1). In contrast, the control configuration recorded 0.74 precision, 0.78 recall, and 0.76 F1.
- Adaptive Learning Effect: Learners following the reinforcement-learning–driven curriculum realized an average 14 % improvement in targeted skill assessments across three iterations, exceeding the 10 % improvement target and demonstrating the planner's ability to accelerate proficiency gains.
- User Engagement & Usability: A/B testing with 30 students showed the experimental interface yielded a 61 % click-through rate (versus 45 % in control), 4.1 minutes average session duration (versus 3.2 minutes), and a System Usability Scale (SUS) score of 76.3 (versus 62). Each of these improvements met or surpassed the 30 % uplift and SUS≥70 benchmarks.
- Instrument Reliability: The SUS questionnaire exhibited Cronbach-α = 0.91 and KMO = 0.88, confirming excellent internal consistency and sampling adequacy.

The empirical evidence demonstrates that integrating real-time labor-market data, a hybrid ML recommendation engine, and an RL-based learning planner within a Kubernetes–orchestrated microservices architecture yields substantial gains over a static, rule-based baseline.

First, the significant reduction in API latency and error rates confirms that containerized microservices—with horizontal autoscaling—can deliver low-latency, reliable recommendations at scale. This responsiveness is critical for user satisfaction in an interactive web setting, as reflected by the high SUS scores.

Second, the use of transformer-based NLP embeddings combined with collaborative filtering proved highly effective in aligning career suggestions with individual profiles. The lift in precision and recall suggests that semantic analysis of job descriptions captures nuanced skill requirements and user interests better than traditional rule sets or shallow classifiers.

Third, the adaptive learning component's 14 % skill improvement highlights the value of reinforcement learning in sequencing educational content. By continuously

incorporating user feedback (quiz results, course completions), the planner tailors pathways that accelerate mastery—an outcome that static learning plans cannot match.

Finally, the marked improvements in click-through rates and session durations underscore the importance of real-time notifications and an engaging dashboard. Timely alerts on emerging job trends and progress milestones appear to sustain user motivation and exploration.

Future Directions.

To build on these findings, future work could:

- Incorporate Sentiment and Contextual Signals: Leverage sentiment analysis on employer reviews or forum discussions to refine career fit scores.
- Expand Localization: Develop Bahasa Indonesia–optimized NLP pipelines to improve recommendation quality for locally phrased job postings.
- Integrate Mentorship Networks: Recommend not only courses but also industry mentors or internship opportunities to bolster practical experience.
- Longitudinal Impact Study: Track real-world career outcomes of users over 6–12 months to evaluate the platform's influence on employability and job satisfaction.

By addressing these avenues, the Smart Career Path Advisor can evolve from a recommendation engine into a comprehensive career ecosystem—bridging the gap between education and employment in Indonesia's dynamic digital economy.

4. CONCLUSION

This study has presented the design, implementation, and evaluation of the Smart Career Path Advisor—an AI-driven, microservice-oriented platform that bridges the growing gap between Indonesian graduates' skills and the dynamic requirements of Industry 4.0. By integrating real-time labor-market scraping, a hybrid ML recommendation engine, reinforcement-learning–based adaptive learning, and a responsive, notification-driven interface—all deployed on a Kubernetes-managed microservices architecture—our system delivers:

- High Performance at Scale. Sub-200 ms API latencies and < 1 % error rates under concurrent user loads validate the platform's robustness and responsiveness.
- Accurate, Contextualized Guidance. Top-5 recommendation precision (0.85), recall (0.89), and F1-score (0.87) exceed industry benchmarks, demonstrating the value of combining collaborative filtering with transformer-based NLP.
- Accelerated Skill Development. A 14 % improvement in targeted assessments confirms the efficacy of the RL-powered learning planner in personalizing study paths.
- Engaging, User-Centric Experience. A 61 % click-through rate, 4.1 min average session, and SUS = 76.3 underscore the platform's ability to capture and sustain learner interest.

Together, these outcomes underscore the importance of a fully integrated, data-driven approach to career guidance—moving beyond static advisement toward continuous, personalized support that adapts to both individual progress and evolving market trends. In doing so, the Smart Career Path Advisor not only validates the underlying thesis—that real-time AI and microservices can revolutionize student career

planning—but also offers a blueprint for institutions seeking to empower their graduates with actionable, up-to-date career intelligence.

Looking ahead, extending this platform with sentiment-based job analytics, deeper localization for Bahasa Indonesia, and real-world mentorship integrations promises to further enhance its impact. We believe that by uniting rigorous AI methods with scalable architecture and user-focused design, the Smart Career Path Advisor can serve as a cornerstone for closing the skills-mismatch gap—not just in Indonesia, but across emerging economies worldwide.

ACKNOWLEDGEMENTS

The authors would like to thank all individuals and institutions who provided support, guidance, and assistance during the completion of this research. Their contributions were highly valuable in the preparation of this paper.

REFERENCES

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 4171–4186). https://doi.org/10.18653/v1/N19-1423

He, K., & Sun, J. (2015). Convolutional neural networks at constrained time cost. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 5353–5360). https://doi.org/10.1109/CVPR.2015.7299156

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778). https://doi.org/10.1109/CVPR.2016.90

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436-444 https://doi.org/10.1038/nature14539

Redmon, J., & Farhadi, A. (2018). YOLOv3: An incremental improvement. *arXiv preprint arXiv:1804.02767*. https://doi.org/10.48550/arXiv.1804.02767

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature, 323*(6088), 533–536. https://doi.org/10.1038/323533a0

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks, 61*, 85–117. https://doi.org/10.1016/j.neunet.2014.09.003

Smith, L. N. (2017). Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)* (pp. 464–472). https://doi.org/10.1109/WACV.2017.58

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems* (Vol. 30, pp. 5998–6008). https://doi.org/10.5555/3295222.3295349