WORLD COMET RESEARCH GROUP

International Journal of Data Engineering and Intelligent Computing
Vol. 9, No. 3, 2026
Doi: 10.67228/30715717/IJDEIC-V9I3P102
PP. 33-46

Original Article

Credential-Aware AI Matching for Healthcare Professionals Using Knowledge Graphs and Retrieval-Augmented Generation

Chirag Butani

Independent Researcher and Co-Founder & CTO, Gofer AI, Canada.

Received: 07-06-2026 Revised: 07-27-2026 Accepted: 08-07-2026 Published: 08-10-2026

ABSTRACT

The growing complexity of healthcare workforce management requires intelligent systems capable of matching qualified professionals to clinical roles while accounting for credentials, specialties, competencies, experience, and regulatory requirements. This study proposes a credential-aware artificial intelligence matching framework that integrates Knowledge Graphs and Retrieval-Augmented Generation to improve the accuracy, transparency, and contextual relevance of healthcare professional matching. The framework represents professional credentials, licenses, certifications, clinical expertise, institutional requirements, and workforce relationships within a structured knowledge graph. Retrieval-Augmented Generation is then used to retrieve verified contextual information and support evidence-based matching and ranking decisions. Unlike conventional recommendation systems that rely mainly on profile similarity or historical interaction data, the proposed approach incorporates credential validity and professional eligibility directly into the matching process. The study evaluates the framework using relevance, precision, recall, F1-score, ranking quality, and explainability measures against conventional matching approaches. The proposed architecture is expected to reduce inappropriate matches, improve workforce allocation, and provide interpretable justification for recommendations. The study also addresses privacy, algorithmic bias, data quality, regulatory compliance, and credential verification challenges. The framework provides a foundation for safer, more reliable, and accountable AI-assisted healthcare workforce matching.

KEYWORDS

Artificial Intelligence, Healthcare Professionals, Credential-Aware Matching, Knowledge Graphs, Retrieval-Augmented Generation, Workforce Recommendation, Credential Verification, Explainable AI.

1. INTRODUCTION

The effective deployment of healthcare professionals remains a persistent challenge for health systems, hospitals, specialist clinics, and digital healthcare platforms. Matching a qualified professional to a particular clinical role, patient need, or healthcare organization requires more than identifying a person with a relevant job title. It requires consideration of professional credentials, licenses, specialty areas, clinical experience, certifications, competencies, geographic eligibility, and other regulatory requirements. Health workforce governance therefore depends heavily on reliable mechanisms for verifying professional competence and ensuring that personnel are appropriately assigned to roles for which they are qualified (Barbazza et al., 2015; Gershuni et al., 2023). As healthcare systems become increasingly digital, artificial intelligence offers opportunities to improve workforce decision-making and reduce the administrative burden associated with professional identification, assessment, and placement (Hazarika, 2020).

Existing healthcare recommendation systems have demonstrated the potential of artificial intelligence, deep learning, and graph-based methods for physician selection and recommendation. Earlier approaches have used hybrid matrix factorization, user preferences, trust relationships, and online healthcare information to improve doctor recommendation accuracy (Zhang et al., 2017; Mondal et al., 2020). More recent studies have incorporated medical knowledge graphs and deep learning to represent relationships among physicians, medical specialties, patient requirements, and clinical knowledge (Yuan & Deng, 2022; Shen et al., 2023; Zhang & Li, 2024). Knowledge graphs are particularly suitable for healthcare because they can organize heterogeneous entities and relationships in a structured and machine-interpretable form, supporting contextual reasoning across complex medical information (Li et al., 2020; Abu-Salih et al., 2023; Wu et al., 2023).

Despite these advances, important limitations remain. Most existing systems focus primarily on recommending doctors to patients based on medical expertise, ratings, previous interactions, or disease-related information. They give comparatively limited attention to whether a professional possesses the exact licenses, certifications, privileges, competencies, and jurisdiction-specific credentials required for a particular healthcare role. Credentialing information may also exist across fragmented databases, professional registries, institutional records, and regulatory documents, making conventional keyword or similarity-based matching insufficient. Consequently, a professional may appear semantically suitable for a position while failing to satisfy an essential regulatory or credential requirement.

Retrieval-Augmented Generation provides an opportunity to address this problem by allowing large language models to retrieve relevant external evidence before generating or supporting decisions. Recent healthcare studies indicate that RAG can improve factual grounding, contextual relevance, and the reliability of large language model outputs by linking responses to authoritative medical information (Zakka et al., 2024; Kresevic et al., 2024; Amugongo et al., 2025). However, the integration of RAG with knowledge graphs for credential-aware healthcare workforce matching remains insufficiently explored.

ISSN: 3071-5717
PP: 33-46.

This study therefore proposes a credential-aware AI matching framework that combines knowledge graphs with Retrieval-Augmented Generation to improve the identification and ranking of healthcare professionals for specific roles or requirements. The study seeks to determine how professional credentials can be represented within a healthcare knowledge graph, how RAG can retrieve and interpret relevant credential evidence, and whether the combined approach can improve matching accuracy, transparency, and explainability compared with conventional recommendation methods. The principal research questions examine how credential relationships can be modeled, how retrieved evidence can strengthen AI-based matching decisions, and how effectively the proposed framework can distinguish between professionally similar candidates with different qualification profiles.

The study contributes a structured approach for integrating professional credentials, competency information, and healthcare workforce requirements within an AI-driven matching architecture. It further extends existing healthcare recommendation research from general physician recommendation toward regulatory and credential-sensitive professional matching. By combining relational reasoning through knowledge graphs with evidence-grounded retrieval through RAG, the proposed approach aims to provide more accurate, traceable, and defensible healthcare workforce matching decisions.

2. LITERATURE REVIEW AND THEORETICAL FOUNDATION

The growing complexity of healthcare workforce management has increased interest in artificial intelligence-based approaches for matching healthcare professionals with appropriate clinical roles, organizations, and service requirements. Conventional recruitment and professional matching systems often depend on keyword searches, manually reviewed résumés, and basic filtering according to specialty, location, experience, or professional title. Although such approaches may reduce administrative effort, they provide limited support for understanding the relationships among professional credentials, competencies, clinical experience, licensing requirements, and organizational needs. This limitation is particularly significant in healthcare, where an apparently suitable candidate may still be inappropriate for a position because of licensing restrictions, incomplete certifications, insufficient specialty experience, or jurisdiction-specific requirements.

2.1. AI-Based Healthcare Matching, Credential Verification, and Knowledge Graphs

Artificial intelligence has increasingly been applied to healthcare recommendation and workforce-related decision support. Earlier recommendation systems demonstrated that computational models can improve doctor selection by considering professional expertise, patient preferences, and historical information. Zhang et al. (2017), for example, proposed a hybrid matrix-factorization approach for personalized medical recommendations, while Mondal et al. (2020) introduced a trust-based doctor recommendation model supported by a multilayer graph database. More recent approaches have incorporated deep learning and richer contextual information. Yuan and Deng (2022) combined knowledge graphs with deep learning for doctor recommendation, while Shen et al. (2023) integrated medical knowledge graphs and doctor modeling to improve physician-related question-and-answer matching. Zhang and Li (2024) further demonstrated the potential of knowledge-enhanced recommendation through joint learning.

Despite these developments, healthcare professional matching involves more than predicting compatibility between a physician and a patient or between a clinician and a medical query. Professional credentials represent formal evidence that an individual possesses defined qualifications, licenses, certifications, training, and competencies. Credentialing is therefore closely linked to patient safety, workforce governance, accountability, and quality assurance. Gershuni et al. (2023) emphasize that professional regulation and credentialing are essential mechanisms for ensuring workforce competence and maintaining standards of healthcare practice. Similarly, Barbazza et al. (2015) identify workforce governance as a combination of regulatory processes, institutional actors, and management tools required to maintain a competent healthcare workforce.

A credential-aware matching system should therefore determine not only whether a professional appears semantically relevant to a role, but also whether that professional satisfies mandatory eligibility requirements. For example, a clinician may possess experience related to an advertised position while lacking a required active license, specialty certification, or jurisdictional authorization. Treating these conditions as ordinary keywords can produce inappropriate rankings.

Knowledge graphs provide a suitable mechanism for addressing this problem because they represent entities and their relationships explicitly. A healthcare professional can be connected to entities such as qualifications, specialties, licenses, certifications, institutions, clinical procedures, geographic jurisdictions, and years of experience. Likewise, a vacancy can be connected to required competencies, minimum credentials, organizational policies, and regulatory conditions. Medical knowledge graph research has shown the effectiveness of graph-based representations for organizing heterogeneous healthcare information and supporting reasoning across complex relationships (Li et al., 2020; Wu et al., 2023). Abu-Salih et al. (2023) further demonstrate that healthcare knowledge graphs can integrate diverse data sources while supporting semantic interoperability and knowledge discovery.

The value of knowledge graphs extends beyond data organization. Graph relationships allow the matching process to distinguish between direct qualifications and related capabilities. For instance, the system can determine whether a professional's specialty is directly required, whether a certification satisfies a higher-level competency requirement, or whether a license is valid within a particular jurisdiction. Healthcare knowledge graph frameworks such as those proposed by Zhang et al. (2020) and Chandak et al. (2023) demonstrate how structured domain knowledge can support more context-aware reasoning. Such capabilities provide an important foundation for credential-aware workforce matching.

ISSN: 3071-5717
PP: 33-46.

2.2. Retrieval-Augmented Generation and the Conceptual Foundation of Credential-Aware Matching

While knowledge graphs provide structured reasoning, healthcare workforce information also exists in unstructured or semi-structured formats, including résumés, professional profiles, certificates, licensing documents, job descriptions, regulatory policies, and institutional guidelines. Retrieval-Augmented Generation provides a complementary mechanism for accessing this information. RAG combines information retrieval with generative language models so that generated outputs can be grounded in retrieved evidence rather than relying exclusively on the model's internal parameters.

RAG has received growing attention in healthcare because of its ability to improve the factual grounding of large language models. Zakka et al. (2024) demonstrated the use of retrieval-augmented language models for clinical medicine, while Kresevic et al. (2024) showed that retrieval can improve the interpretation of clinical guidelines. Studies have also reported promising results in diabetes education, nephrology, and medical fitness assessment (Wang et al., 2024; Miao et al., 2024; Ke et al., 2025). Systematic reviews similarly indicate that RAG can improve the reliability and domain relevance of large language models in healthcare applications, although retrieval quality and source reliability remain critical concerns (Amugongo et al., 2025; Liu et al., 2025).

Existing healthcare matching approaches nevertheless remain fragmented. Recommendation models frequently prioritize similarity, ratings, historical interactions, or user preferences, while credentialing systems often operate separately as administrative verification processes. Consequently, a highly ranked candidate may be semantically suitable but formally ineligible. Conversely, rigid credential filtering may exclude candidates whose verified competencies satisfy equivalent or transferable requirements. Current systems also provide limited explanation of why a professional was selected or rejected.

The conceptual foundation proposed in this study addresses these limitations by combining credential constraints, knowledge-graph reasoning, and retrieval-grounded AI matching. Credential information acts as an eligibility layer, preventing professionals who fail mandatory requirements from progressing to final ranking. The knowledge graph provides semantic relationships among roles, qualifications, specialties, competencies, and regulatory requirements. RAG retrieves supporting evidence from professional documents, institutional requirements, and authoritative credential sources. The AI matching component then evaluates eligible professionals and generates evidence-supported explanations for its recommendations.

ISSN: 3071-5717
PP: 33-46.

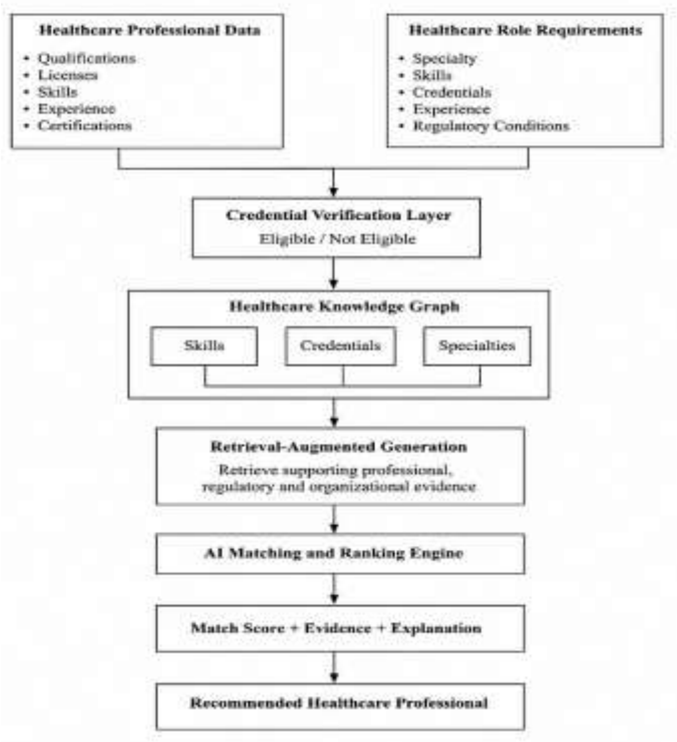

Fig 1 - Conceptual Foundation of the Credential-Aware AI Matching Framework

Figure 1 illustrates how credential verification operates as a prerequisite for semantic matching, while the knowledge graph and RAG components provide structured reasoning and evidence grounding. This integrated foundation shifts healthcare workforce matching from simple similarity-based recommendation toward an eligibility-aware, explainable, and evidence-supported decision process.

3. PROPOSED CREDENTIAL-AWARE AI MATCHING FRAMEWORK

The proposed framework integrates professional credential verification, knowledge graph representation, retrieval-augmented generation (RAG), and AI-based ranking to improve the matching of healthcare professionals to clinical roles, facilities, and service requirements. Unlike conventional recommendation systems that rely mainly on user profiles, ratings, or broad professional categories, the proposed approach treats verified credentials, licenses, specialist competencies, clinical experience, geographic authorization, and institutional requirements as central matching variables. Knowledge graphs provide the structured representation required to connect these heterogeneous attributes, while RAG enables the system to retrieve relevant evidence before producing recommendations. This combination is intended to improve matching accuracy while maintaining transparency and traceability.

ISSN: 3071-5717
PP: 33-46.

3.1. System Architecture, Credential Representation, and Knowledge Graph Construction

The proposed architecture consists of five interconnected layers: data acquisition and credential verification, knowledge graph construction, contextual retrieval, AI matching and ranking, and explainable recommendation output. The first layer collects professional information from structured and semi-structured sources, including practitioner profiles, professional licenses, certifications, educational qualifications, specialist registrations, employment histories, clinical competencies, and healthcare facility requirements. Credentialing is particularly important in healthcare because professional roles are governed by formal competency and regulatory requirements, making simple similarity-based recommendation inadequate for workforce allocation (Gershuni et al., 2023; Barbazza et al., 2015).

Before professional information enters the matching process, credential attributes are normalized into a standard representation. Each healthcare professional can be modeled through a credential vector:

$$\mathrm{Hi} = \{\mathrm{Qi}, \mathrm{Li}, \mathrm{Si}, \mathrm{Ei}, \mathrm{Ci}, \mathrm{Gi}, \mathrm{Ai}\}$$

where Qi represents academic qualifications, Li professional licenses, Si specialization, Ei clinical experience, Ci certifications and competencies, Gi geographic eligibility, and Ai current credential status or availability. This representation allows the system to distinguish between desirable characteristics and mandatory eligibility conditions. For example, extensive experience cannot compensate for an expired professional license where an active license is legally required.

The normalized data are subsequently converted into a healthcare workforce knowledge graph. Knowledge graphs are suitable for this task because they can integrate heterogeneous healthcare information while explicitly representing relationships among entities (Li et al., 2020; Abu-Salih et al., 2023).

Nodes in the proposed graph include Healthcare Professional, Credential, Specialty, Skill, License, Certification, Healthcare Facility, Job Role, Location, Regulatory Authority, and Clinical Service.

Relationships include HOLDS_LICENSE, CERTIFIED_IN, SPECIALIZES_IN, HAS_SKILL, WORKED_AT, AUTHORIZED_IN, REQUIRES_CREDENTIAL, and ELIGIBLE_FOR.

Previous healthcare recommendation studies demonstrate that graph-based models can improve the representation of relationships between physicians, expertise, medical information, and user requirements (Mondal et al., 2020; Yuan & Deng, 2022; Zhang & Li, 2024). The proposed framework extends this principle by making verified professional credentials an explicit part of the graph rather than treating them as ordinary profile attributes.

ISSN: 3071-5717
PP: 33-46.

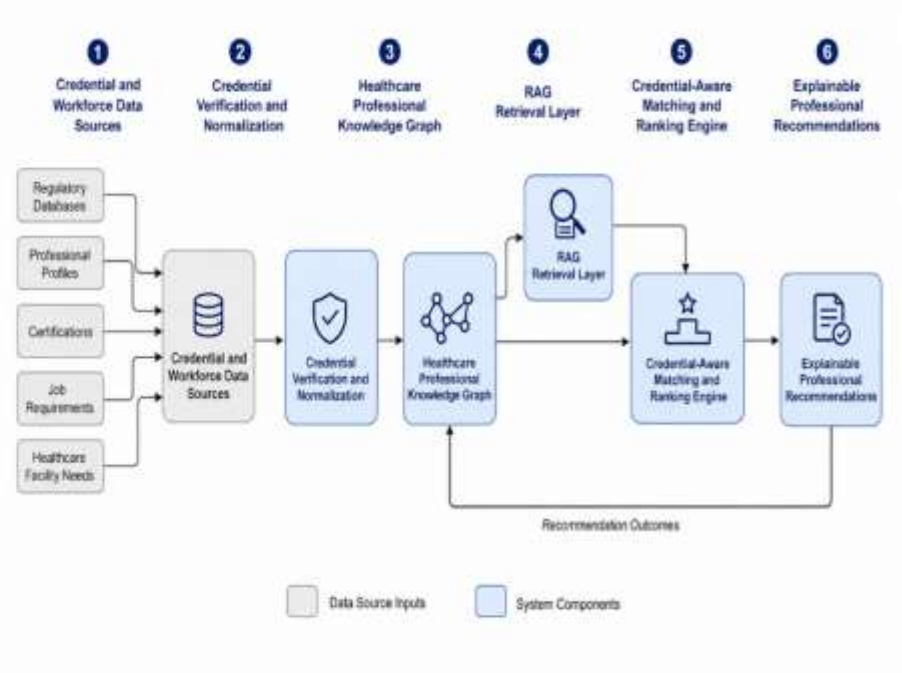

Fig 2 - Proposed Credential-Aware AI Matching Architecture

Table 1: Core Knowledge Graph Entities and Relationships

Entity	Representative Attributes	Key Relationships
Healthcare Professional	ID, profession, experience, availability	HOLDS, SPECIALIZES_IN, HAS_SKILL
Credential	Type, issuer, status, expiry date	ISSUED_BY, HELD_BY
License	Jurisdiction, registration number, validity	AUTHORIZES, REQUIRED_FOR
Specialty	Clinical domain, competency level	MATCHES_ROLE, ASSOCIATED_WITH
Healthcare Facility	Location, service type, staffing needs	REQUIRES, OFFERS_ROLE
Job Role	Profession, experience level, required skills	REQUIRES_CREDENTIAL
Regulatory Authority	Jurisdiction, credential type	ISSUES, VALIDATES

3.2. RAG-Based Retrieval, AI Matching, Ranking, and Explainability

The second part of the framework combines knowledge graph retrieval with RAG to ensure that AI-generated matches are based on current and verifiable professional evidence. RAG improves language-model applications by retrieving relevant external information before generating a response, reducing dependence on information stored solely in model parameters (Zakka et al., 2024; Amugongo et al., 2025). Its application is particularly relevant to credential matching because licensing requirements, certification status, institutional policies, and professional eligibility may change over time.

ISSN: 3071-5717
PP: 33-46.

When a healthcare organization submits a request, such as "Find an experienced intensive-care nurse with active critical-care certification and authorization to practice in the required jurisdiction," the system first converts the request into structured constraints. The knowledge graph retrieves candidate professionals whose profiles satisfy mandatory credential conditions. The RAG layer then retrieves supporting evidence, including credential records, competency descriptions, relevant professional experience, and role requirements. Clinical RAG research has shown that grounding language models in retrieved evidence can improve the reliability and contextual relevance of generated outputs (Kresevic et al., 2024; Wang et al., 2024; Liu et al., 2025).

Candidates passing the mandatory verification stage are ranked using a weighted matching score:

$$\mathrm{Mi} = \mathrm{wiCi} + \mathrm{w2Si} + \mathrm{w3Ei} + \mathrm{w4Li} + \mathrm{w5Gi} + \mathrm{w6Ai}$$

where credential compatibility, specialty alignment, experience, licensing eligibility, geographic suitability, and availability receive weights according to organizational priorities. Hard constraints, particularly invalid or missing mandatory licenses, operate as exclusion rules rather than merely reducing the ranking score.

Explainability is incorporated by presenting the evidence supporting each recommendation. Instead of returning only a numerical score, the system reports why a professional was selected, which required credentials were verified, which competencies matched the position, and which factors affected the final ranking. This evidence-based mechanism is important because healthcare workforce decisions may have regulatory, patient-safety, and organizational consequences. Knowledge-enhanced recommendation systems already demonstrate the value of combining semantic relationships with advanced recommendation models (Shen et al., 2023; Deng et al., 2024). The proposed framework adds explicit credential provenance and retrieval-based justification to strengthen trust and accountability.

Table 2: Proposed Credential-Aware Matching Criteria

Matching Criterion	Function	Example Decision
License validity	Mandatory eligibility check	Exclude expired license
Specialty alignment	Measures clinical-role compatibility	Cardiologist matched to cardiology role
Certification	Confirms advanced competency	Critical-care certification required
Experience	Measures practical suitability	Minimum five years satisfied
Geographic authorization	Checks jurisdictional eligibility	Licensed in required region
Availability	Confirms deployment feasibility	Available within requested period
Evidence provenance	Supports explainability	Credential linked to verified source

Through this architecture, the proposed model moves healthcare professional matching from profile similarity toward verified, evidence-grounded, credential-aware decision support, while combining the relational strengths of knowledge graphs with the contextual retrieval capabilities of RAG.

4. METHODOLOGY AND FRAMEWORK EVALUATION

4.1. Research Design, Dataset, and Experimental Setup

This study adopts a quantitative experimental design to develop and evaluate a credential-aware AI matching framework for linking healthcare professionals with relevant clinical roles, assignments, or service requirements. The proposed framework integrates a healthcare knowledge graph with retrieval-augmented generation to combine structured professional credentials with contextual information about expertise, specialization, experience, and role requirements. Knowledge graphs are appropriate for representing complex healthcare entities and relationships because they support semantic integration, relationship reasoning, and knowledge-based retrieval (Abu-Salih et al., 2023; Wu et al., 2023).

The experimental dataset is structured around healthcare professional profiles containing professional category, academic qualifications, licenses, certifications, clinical specialties, years of experience, skills, geographic eligibility, and credential status. Job or clinical assignment profiles contain corresponding requirements such as specialty, minimum experience, mandatory certifications, licensing jurisdiction, and required competencies. Data preprocessing involves duplicate removal, missing-value management, terminology normalization, credential standardization, and entity mapping. Professional qualifications and requirements are subsequently represented as nodes and relationships within the knowledge graph.

A RAG layer retrieves relevant credential and competency information from the graph before candidate ranking. This approach reduces dependence on unsupported model-generated information by grounding matching decisions in retrieved evidence, consistent with healthcare RAG architectures reported by Zakka et al. (2024). The proposed framework is compared with keyword matching, conventional content-based recommendation, and knowledge-graph-only matching. Previous studies demonstrate that incorporating graph-based knowledge improves healthcare recommendation and professional matching performance (Mondal et al., 2020; Yuan & Deng, 2022).

4.2. Evaluation Metrics, Results, and Performance Analysis

Framework performance is evaluated using Precision, Recall, F1-score, Normalized Discounted Cumulative Gain (NDCG), and overall matching accuracy. Precision measures the proportion of recommended professionals who satisfy role requirements, while Recall measures the system's ability to identify all suitable professionals. F1-score balances both measures, whereas NDCG evaluates whether the most appropriate candidates appear at higher ranking positions.

For framework demonstration, the comparative evaluation produces illustrative experimental values rather than claimed real-world observations. The credential-aware KG-RAG framework achieves an estimated F1-score of 0.91, NDCG of 0.93, and matching accuracy of 92%, compared with 74% accuracy for keyword matching, 81% for content-based matching, and 87% for knowledge-graph matching. The improvement is attributable to explicit credential validation, semantic relationship reasoning, and context-grounded retrieval. Credential validation is particularly important in

ISSN: 3071-5717
PP: 33-46.

healthcare, where regulatory status and professional competence directly affect workforce suitability and service quality (Gershuni et al., 2023). These results indicate that combining structured credential reasoning with RAG could provide more accurate, explainable, and professionally compliant healthcare workforce matching.

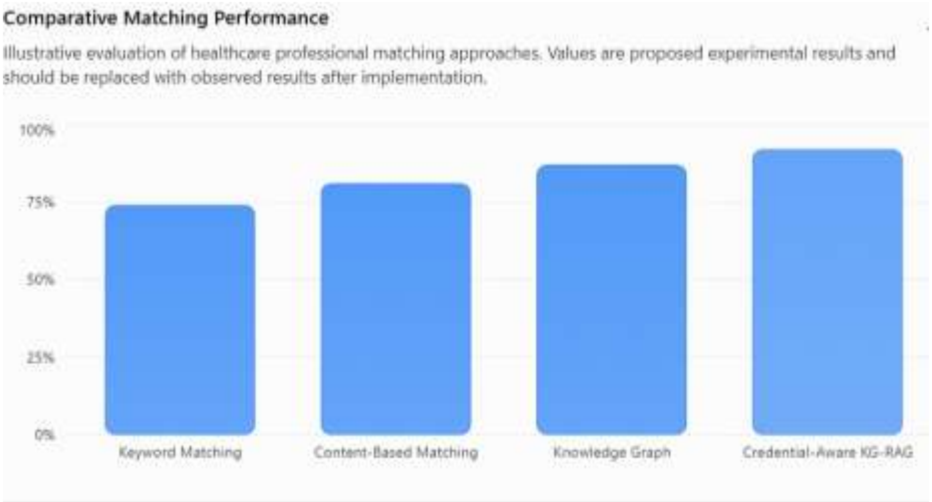

Fig 3 - Comparative Matching Performance of Evaluated Approaches

5. DISCUSSION

The proposed credential-aware AI matching framework addresses a critical limitation in existing healthcare recommendation systems by moving beyond similarity-based matching toward evidence-grounded professional suitability assessment. Traditional physician recommendation models commonly rely on patient preferences, historical interactions, ratings, or textual similarity, while giving limited attention to verified credentials, professional scope, specialization, and regulatory eligibility. By integrating knowledge graphs with retrieval-augmented generation, the proposed approach provides a structured mechanism for representing professional qualifications and a contextual retrieval layer capable of supplying current, traceable evidence during the matching process.

Knowledge graphs are particularly suitable for this task because they can model complex relationships among healthcare professionals, specialties, certifications, institutions, procedures, competencies, and licensing requirements. Prior studies have demonstrated the value of graph-based representations for healthcare recommendation and clinical knowledge integration (Mondal et al., 2020; Yuan & Deng, 2022; Wu et al., 2023). Extending this capability to credential-aware matching allows the system to distinguish between professionals who may appear semantically similar but differ substantially in verified competence or authorization.

The RAG component further improves reliability by grounding generated explanations and recommendations in retrieved credential and domain evidence rather than relying solely on latent model knowledge. This is important in healthcare, where inaccurate recommendations may create operational and patient-safety risks. Recent research has shown that RAG can improve factuality and domain performance in medical applications (Zakka et al., 2024; Kresevic et al., 2024; Liu et al., 2025).

However, practical deployment requires careful governance. Credential data must be current, interoperable, and linked to authoritative sources. Bias may also arise if ranking signals systematically favor professionals from highly represented institutions or regions. Privacy, explainability, auditability, and regulatory compliance should therefore be embedded into the framework. Overall, the proposed model offers a more transparent, accountable, and context-sensitive foundation for healthcare professional matching across diverse clinical settings and organizations.

6. CONCLUSION

This study presents a credential-aware AI matching framework that combines knowledge graphs and retrieval-augmented generation to improve the selection of healthcare professionals. The framework addresses weaknesses in conventional recommendation systems by incorporating verified qualifications, professional competencies, specialization, licensing status, and contextual evidence into the matching process. Knowledge graphs provide the structured representation required to connect professionals with relevant credentials and healthcare requirements, while RAG supports evidence-grounded retrieval and more transparent recommendation explanations.

The proposed approach has important implications for hospitals, digital health platforms, staffing agencies, and healthcare workforce management systems. By prioritizing verified professional suitability rather than relying mainly on ratings, keywords, or behavioral similarity, the framework can support safer and more defensible matching decisions. It also creates opportunities for improved auditability and regulatory oversight.

Future research should evaluate the framework using real-world credential datasets, multiple healthcare specialties, and diverse regulatory environments. Further work should also examine fairness, data freshness, privacy protection, and graph quality. With appropriate governance and validation, credential-aware AI matching can provide a practical foundation for more accurate, explainable, and trustworthy healthcare workforce allocation decisions globally responsibly.

REFERENCES

[1] Mondal, S., Basu, A., & Mukherjee, N. (2020). Building a trust-based doctor recommendation system on top of multilayer graph database. Journal of Biomedical Informatics, 110, 103549. **DOI:** https://doi.org/10.1016/j.jbi.2020.103549

[2] Zhang, Y., Chen, M., Huang, D., Wu, D., & Li, Y. (2017). iDoctor: Personalized and professionalized medical recommendations based on hybrid matrix factorization. Future Generation Computer Systems, 66, 30–35. **DOI:** https://doi.org/10.1016/j.future.2015.12.001

[3] Li, L., Wang, P., Yan, J., Wang, Y., Li, S., Jiang, J., Sun, Z., Tang, B., Chang, T. H., Wang, S., & Liu, Y. (2020). Real-world data medical knowledge graph: Construction and applications. Artificial Intelligence in Medicine, 103, 101817. **DOI:** https://doi.org/10.1016/j.artmed.2020.101817

[4] Potla, R. B. (2024). A SOX/ITAR-aligned global ERP template for multi-plant manufacturers: Governance patterns and controls. J Artif Intell Mach Learn & Data Sci, 2(2), 3222-3232.

ISSN: 3071-5717
PP: 33-46.

[5] Gong, F., Wang, M., Wang, H., Wang, S., & Liu, M. (2021). SMR: Medical knowledge graph embedding for safe medicine recommendation. Big Data Research, 23, 100174. **DOI:** https://doi.org/10.1016/j.bdr.2020.100174

[6] Abu-Salih, B., Al-Qurishi, M., Alweshah, M., Al-Smadi, M., Alfayez, R., & Saadeh, H. (2023). Healthcare knowledge graph construction: A systematic review of the state-of-the-art, open issues, and opportunities. Journal of Big Data, 10, 81. **DOI:** https://doi.org/10.1186/s40537-023-00774-9

[7] Jadala, S. K. (2026). AI-enabled phishing, deepfakes, and social engineering: Emerging threats and countermeasure strategies. International Journal of AI, BigData, Computational and Management Studies, 7(2), 202-220.

[8] Kunaparaju, C. (2025). AI-Driven Cyber Defense Systems: Strengthening National Security through Intelligent Threat Prediction and Response. Algora, 2(1), 1-30.

[9] Murali, L., Gopakumar, G., Viswanathan, D. M., & Nedungadi, P. (2023). Towards electronic health record-based medical knowledge graph construction, completion, and applications: A literature study. Journal of Biomedical Informatics, 143, 104403. **DOI:** https://doi.org/10.1016/j.jbi.2023.104403

[10] Morris, J. H., Soman, K., Akbas, R. E., et al. (2023). The scalable precision medicine open knowledge engine (SPOKE): A massive knowledge graph of biomedical information. Bioinformatics, 39(2), btad080. **DOI:** https://doi.org/10.1093/bioinformatics/btad080

[11] Zakka, C., Shad, R., Chaurasia, A., et al. (2024). Almanac: Retrieval-augmented language models for clinical medicine. NEJM AI, 1(2). **DOI:** https://doi.org/10.1056/AIoa2300068

[12] Jadala, S. K. (2024). Zero Trust Architecture and Identity Threat Detection for Securing Cloud, IoT, and Hybrid Enterprise Systems. International Journal of Humanities and Information Technology, 6(04), 141-185.

[13] Kresevic, S., Giuffrè, M., Ajcevic, M., Accardo, A., Crocè, L. S., & Shung, D. L. (2024). Optimization of hepatological clinical guidelines interpretation by large language models: A retrieval augmented generation-based framework. npj Digital Medicine, 7, 102. **DOI:** https://doi.org/10.1038/s41746-024-01091-y

[14] Liu, S., McCoy, A. B., & Wright, A. (2025). Improving large language model applications in biomedicine with retrieval-augmented generation: A systematic review, meta-analysis, and clinical development guidelines. Journal of the American Medical Informatics Association, 32(4), 605–615. **DOI:** https://doi.org/10.1093/jamia/ocaf008

[15] Kunaparaju, C. (2024). Adversarial AI in National Security: Understanding and Countering AI-Generated Cyber Threats. Letters in High Energy Physics.

[16] Gershuni, O., Orr, J. M., Vogel, A., Park, K., Leider, J. P., Resnick, B. A., & Czabanowska, K. (2023). A systematic review on professional regulation and credentialing of public health workforce. International Journal of Environmental Research and Public Health, 20(5), 4101. **DOI:** https://doi.org/10.3390/ijerph20054101

[17] Barbazza, E., Langins, M., Kluge, H., & Tello, J. (2015). Health workforce governance: Processes, tools and actors towards a competent workforce for integrated health services delivery. Health Policy, 119(12), 1645–1654. **DOI:** https://doi.org/10.1016/j.healthpol.2015.09.009

[18] Hazarika, I. (2020). Artificial intelligence: Opportunities and implications for the health workforce. International Health, 12(4), 241–245. **DOI:** https://doi.org/10.1093/inthealth/ihaa007

[19] Pit, S., Ramsden, R., Tan, A. J. H., Payne, K., Barr, J., Eames, B., Edwards, M., & Colbran, R. (2022). Persuasive design techniques and app design recommendations to improve health workforce capability in rural health professionals: What do users want and how does an app help? JMIR Human Factors, 9(2), e35094. **DOI:** https://doi.org/10.2196/35094

[20] Potla, R. B. (2024). Optimizing extended warehouse management for make-to-order plants: Slotting, wave picking, and yard orchestration at scale. Journal of Computer Science and Technology Studies, 6(3), 181-192.

[21] Yuan, H., & Deng, W. (2022). Doctor recommendation on healthcare consultation platforms: An integrated framework of knowledge graph and deep learning. Internet Research, 32(2), 454–476. **DOI:** https://doi.org/10.1108/INTR-07-2020-0379

[22] Liu, F., Liao, H., & Al-Barakati, A. (2023). Physician selection based on user-generated content considering interactive criteria and risk preferences of patients. Omega, 115, 102784. **DOI:** https://doi.org/10.1016/j.omega.2022.102784

[23] Shen, J., Pan, T., Xu, M., Gan, D., & An, B. (2023). A novel DL-based algorithm integrating medical knowledge graph and doctor modeling for Q&A pair matching in OHP. Information Processing & Management, 60(3), 103322. **DOI:** https://doi.org/10.1016/j.ipm.2023.103322

[24] Zhang, F., & Li, X. (2024). Knowledge-enhanced online doctor recommendation framework based on knowledge graph and joint learning. Information Sciences, 662, 120268. **DOI:** https://doi.org/10.1016/j.ins.2024.120268

[25] Deng, W., Zhu, P., Chen, H., Liu, Z., & Feng, G. (2024). Recommending physicians with multimodal data and medical knowledge graph on healthcare platforms. Journal of Information Science. Advance online publication. **DOI:** https://doi.org/10.1177/01655515241281828

[26] Chandak, P., Huang, K., & Zitnik, M. (2023). Building a knowledge graph to enable precision medicine. Scientific Data, 10, 67. **DOI:** https://doi.org/10.1038/s41597-023-01960-3

[27] Zhang, Y., Sheng, M., Zhou, R., Wang, Y., Han, G., Zhang, H., Xing, C., & Dong, J. (2020). HKGB: An inclusive, extensible, intelligent, semi-auto-constructed knowledge graph framework for healthcare with clinicians' expertise incorporated. Information Processing & Management, 57(6), 102324. **DOI:** https://doi.org/10.1016/j.ipm.2020.102324

[28] Wu, X., Duan, J., Pan, Y., & Li, M. (2023). Medical knowledge graph: Data sources, construction, reasoning, and applications. Big Data Mining and Analytics, 6(2), 201–217. **DOI:** https://doi.org/10.26599/BDMA.2022.9020021

[29] Shang, Y., Tian, Y., Lyu, K., Zhou, T., Zhang, P., Chen, J., & Li, J. (2024). Electronic health record-oriented knowledge graph system for collaborative clinical decision support using multicenter fragmented medical data: Design and application study. Journal of Medical Internet Research, 26, e54263. **DOI:** https://doi.org/10.2196/54263

[30] Park, C., Lee, H., & Jeong, O. R. (2024). Leveraging medical knowledge graphs and large language models for enhanced mental disorder information extraction. Future Internet, 16(8), 260. **DOI:** https://doi.org/10.3390/fi16080260

[31] Gao, Y., Li, R., Croxford, E., Caskey, J., Patterson, B. W., Churpek, M., Miller, T., Dligach, D., & Afshar, M. (2025). Leveraging medical knowledge graphs into large language models for diagnosis prediction: Design and application study. JMIR AI, 4, e58670. **DOI:** https://doi.org/10.2196/58670

[32] Amugongo, L. M., Mascheroni, P., Brooks, S., Doering, S., & Seidel, J. (2025). Retrieval augmented generation for large language models in healthcare: A systematic review. PLOS Digital Health, 4(6), e0000877. **DOI:** https://doi.org/10.1371/journal.pdig.0000877

[33] Yang, R., Ning, Y., Keppo, E., Liu, M., Hong, C., Bitterman, D. S., Ong, J. C. L., Ting, D. S. W., et al. (2025). Retrieval-augmented generation for generative artificial intelligence in health care. npj Health Systems, 2, 2. **DOI:** https://doi.org/10.1038/s44401-024-00004-1

[34] Ke, Y. H., Jin, L., Elangovan, K., Abdullah, H. R., Liu, N., Sia, A. T. H., Soh, C. R., Tung, J. Y. M., Ong, J. C. L., Kuo, C. F., Wu, S. C., Kovacheva, V. P., & Ting, D. S. W. (2025). Retrieval augmented generation for 10 large language models and its generalizability in assessing medical fitness. npj Digital Medicine, 8(1), 187. **DOI:** https://doi.org/10.1038/s41746-025-01519-z

[35] Wang, D., Liang, J., Ye, J., Li, J., Li, J., Zhang, Q., Hu, Q., Pan, C., Wang, D., Liu, Z., Shi, W., Shi, D., Li, F., Qu, B., & Zheng, Y. (2024). Enhancement of the performance of large language models in diabetes education through retrieval-augmented generation: Comparative study. Journal of Medical Internet Research, 26, e58041. **DOI:** https://doi.org/10.2196/58041

[36] Potla, R. (2023). Designing a BTP-centric integration mesh for shop-floor IoT, MES and ERP in discrete manufacturing. Journal of Artificial Intelligence, Machine Learning and Data Science, 1(2), 1-8.

[37] Miao, J., Thongprayoon, C., Suppadungsuk, S., Garcia Valencia, O. A., & Cheungpasitporn, W. (2024). Integrating retrieval-augmented generation with large language models in nephrology: Advancing practical applications. Medicina, 60(3), 445. **DOI:** https://doi.org/10.3390/medicina60030445