

Retrieval-Augmented Generation Architectures for Longitudinal Pediatric EHR Analysis

Vinod Rufus Motani

AI/ML Engineer, Independent Researcher, USA

motanivinodrufus18@gmail.com

ORCID: 0009-0005-3459-0859

Abstract

This paper evaluates Retrieval-Augmented Generation models for longitudinal electronic health records of pediatrics. Pediatric electronic health records include rich clinical information about a patient throughout several healthcare visits. The language models tend to overlook historical contexts while conducting clinical analysis processes. However, the RAG integrates information retrieval with Large Language Models and provides grounded clinical responses. In such a manner, the RAG retrieves the pertinent medical records of the patients before generating the evidence-based clinical summaries and recommendations. The present study employed qualitative methodology that involved secondary data from the healthcare literature and datasets. The inductive strategy was used in the systematic interpretation of the published clinical evidence. The results indicated an improvement in the retrieval of longitudinal patient histories during multiple visits. The RAG framework incorporated diagnoses, medications, lab results, immunizations, and clinical notes. The retrieved information helped to provide effective clinical decision support using evidence-based reasoning and output. The analysis demonstrated better continuity of care using integrated patient timeline records. Privacy was still important since the records of children contain confidential developmental and family history details. RAG ensured that unnecessary exposure was minimized by only extracting clinical information needed for the analysis. Age-related interpretations of growth, developmental, milestones, medications, and laboratory value were some of the key findings. Retrieval ensured that time spent on reviewing pediatric clinical documentation was minimized while providing care. The findings have shown that RAG ensures accurate privacy-preserving and efficient analysis of pediatric EHR.

Keywords: Retrieval-Augmented Generation, pediatric electronic health records, longitudinal patient history, clinical decision support, healthcare artificial intelligence, thematic analysis.

Introduction

Longitudinal pediatric Electronic Health Records have very comprehensive details of all the medical care experiences of that particular child. The data includes diagnoses, prescriptions, lab tests, procedures, and clinical documentation from various encounters of the patient. It is challenging to analyze the information contained in longitudinal pediatric EHR because the dataset is quite vast and constantly changing. There are many significant clinical insights that can get lost in the documentation provided. The Retrieval Augmented Generation (RAG) architecture makes it easier to analyze the information in the EHRs. RAG can be helpful in comprehending patient histories at various points in time throughout the process (Shang *et al.*, 2021). The practice of pediatric medicine needs the knowledge of disease development, responses to therapy, and results after treatment. In addition, RAG helps healthcare professionals in accessing all the important medical information from patient histories that is necessary for answering each specific clinical query. Thus, RAG technology can be used to improve longitudinal pediatric EHR analysis.

Methods

Lewis *et al.*, (2020) states that a qualitative research methodology since it looked at the role of the Retrieval-Augmented Generation (RAG) architecture in longitudinal pediatric electronic health record (EHR) analysis from the available literature and clinical practices. A qualitative approach is ideal for research that involves looking at conceptual issues, clinical processes and information search and retrieval practices and not numerical measurements.




Figure 1: Qualitative Research Framework for RAG-Based Pediatric EHR Analysis

(Source: Self-created)

The research involved the use of secondary data collection methods where information was sought from peer reviewed journal articles, healthcare research studies, publicly available clinical data and scholarly literature on RAG, Large Language Models and pediatric EHR analysis. The use of secondary data ensured that evidence was gathered while excluding any confidential patient records (Oelsner *et al.*, 2020). Inductive reasoning was utilized in this research study where findings were drawn from already existing literature and not from pre-existing hypotheses. The gathered literature was analyzed through thematic analysis in order to determine recurring themes found in the selected research. Themes such as the retrieval of clinical information, patient timeline inclusion, contextual reasoning, longitudinal care continuity, and privacy-conscious healthcare systems were considered during the analysis (Gumiel *et al.*, 2021). Similar results were sorted into common themes and interpreted in order to understand how RAG systems enhance pediatric EHR longitudinal analysis. This approach allowed for a systematic review of existing literature and revealed the contribution of RAG systems in the pediatric clinical information analysis process.

Results

Retrieval of Relevant Clinical Information from Longitudinal Pediatric EHRs

The findings revealed that RAG models enhance the retrieval of clinically relevant information in longitudinal pediatric Electronic Health Records (EHRs) (Vest *et al.*, 2020). In RAG models, patients' records are retrieved first and used in generating clinical responses. It helps the model to leverage historical data through repeated healthcare interactions. In addition, it helps in the accurate interpretation of diagnoses, medications, laboratory tests, immunizations, growth records, and clinical notes throughout the care of children.

Feature	Data Source	Process	Outcome	Benefit
Historical Data	Past Visits	Record Search	Context	Continuity
Medication History	Prescriptions	Information Retrieval	Accuracy	Safety
Growth Records	EHR Data	Timeline Review	Monitoring	Tracking
Lab Results	Clinical Tests	Data Extraction	Insights	Diagnosis
Follow-Up Notes	Prior Visits	Record Linking	Summary	Efficiency
Longitudinal EHRs	Multi-Visit Data	RAG Retrieval	Full Profile	Care Quality

Table 1: Retrieval of Relevant Clinical Information from Longitudinal Pediatric EHRs

The research revealed some of the challenges of retrieval from pediatric EHRs. Relevant clinical information is often found in unstructured notes, family histories, and treatment stories, hence making retrieval very difficult. Furthermore, pediatric EHRs have less content density compared to adult EHRs and are scattered in different hospitals, hence making retrieval challenging (Pai *et al.*, 2021). RAG models overcame this challenge through the retrieval of contextually relevant documents. Research in large databases of pediatrics demonstrates the significance of longitudinal retrieval. PEDSnet database included 9,107,649 people, out of which 7,852,081 were children under 21 years old from 2010 to 2023. Out of those 7,852,081, 1,552,726 (19.8%) suffered from mental health problems or symptoms. Those patients also had longer follow-up durations at 5.1 years as compared to the others who had a follow-up period of 2.4 years. This evidence clearly indicates that longitudinal retrieval is necessary in pediatric data.

Contextual Understanding of Patient History Through RAG Architectures

It is evident that retrieval-augmented generation (RAG) architecture enhances contextual comprehension through retrieval of relevant historical data prior to the generation of clinical responses. Contrary to large language models (LLMs), RAG ensures that the generated responses are based on verified clinical records to mitigate hallucination and enhance the interpretation of pediatric patient's history over time (Rosset *et al.*, 2020). The combination of retrieval from medical knowledge bases and language generation makes sure that diagnoses, medications, lab reports, and past medical encounters are correlated through the RAG architecture.




Figure 2: Simple LLM Query Model

(Source: Pureinsights, 2023)

In this case, there were four essential contextual abilities of the RAG architecture in pediatric EHRs. First, RAG was able to retrieve the patient's diagnosis and treatment history. Secondly, it integrated structured and unstructured clinical data (Sinhababu *et al.*, 2021). Third, it generated evidence-based responses through citation of clinical guidelines and established medical literature. Fourth, it developed transparent clinical summaries which articulated the rationale for recommendations. The longitudinal EHR research has shown the relevance of a multi-visit context. For instance, during an evaluation in the MIMIC-IV dataset, there were patients suffering from chronic ailments and with multiple visits to hospitals, such as a patient who had nine hospital visits and 30 discharge summaries and radiology reports. The graph-based preprocessing framework integrated diagnosis, procedures, medications, laboratory events, radiology reports, and discharge notes into patient histories, making possible the holistic extraction from all the visits.

Integration of Multi-Visit Clinical Records for Continuity of Care

From the results, it is clear that the incorporation of multi-visit clinical records using Retrieval-Augmented Generation (RAG) leads to better continuity of care since it integrates the patient records across multiple visits (Hewa *et al.*, 2020). RAG uses previous records in different visits to generate a single clinical timeline, allowing clinicians to track the progress of the diseases, reaction to treatments, medication histories, and results of follow-ups without going through the previous records. One of the studies done on healthcare in Istanbul has shown the significance of incorporating multiple visits into the records of the patient. This study had 11,101 index visits from 225 family medicine units with 26,743 healthcare visits over a 21-day follow-up. The mean number of visits made per patient was 2.41.

Aspect	Study Data	Value	Timeline	RAG Benefit
Location	Istanbul	Healthcare	Multi-Visit	Continuity
Index Visits	11,101	Patients	Follow-Up	Timeline
Family Units	225	Records	Integrated	Retrieval
Total Visits	26,743	Mean 2.41	Per Patient	Tracking
First Period	3,541 Repeats	48.2%	Days 1–7	History
Second Period	61.2% Visits	Secondary Care	Days 8–21	Full Profile

Table 2: Impact of Multi-Visit Clinical Records on Continuity of Care Using RAG

The frequency of visitations brought out the importance of continuity even more. In the first 7 days, there were 3,541 recurring visits to family physicians which represented 48.2% of all recurring visits made at the primary level. Visits to secondary and tertiary care facilities continued to be high throughout the second period of follow-up with 61.2% taking place in days 8 to 21. RAG models aid in ensuring continuity by enabling retrieval of information from the series of visits and giving a full patient profile (Hashem *et al.*, 2020).

Improvement in Clinical Decision Support and Diagnostic Reasoning

The results show that Retrieval Augmented Generation (RAG) enhances clinical decision-making support through the integration of patient-specific information from EHRs and evidence-based medical knowledge. RAG is capable of retrieving all relevant information such as historical records, clinical guidelines, lab reports, and treatment information before making recommendations, thus enhancing diagnostic reasoning and avoiding unnecessary clinical responses (Zhu *et al.*, 2021).




Figure 3: Three-stage intelligent support of clinical decision making for higher trust, validity, and explainability

(Source: Kovalchuk *et al.*, 2022)

Through the use of RAG, clinicians will be able to make sense of the complex pediatric records in one clinical process. There were various improvements in diagnostic reasoning observed. Structured information from EHRs was combined with unstructured information in clinical notes in order to obtain a patient's full history (Zhang *et al.*, 2020). The market data also reflects the increasing influence of AI-driven Clinical Decision Support Systems (CDSS). The world CDSS market is estimated to grow from USD 2.72 billion in 2025 to USD 4.94 billion in 2031, with a CAGR of 10.49%. Knowledge-based CDSS contributed 61.56% to the overall market revenue in 2025, whereas the cloud-based CDSS is

projected to grow at a CAGR of 16.85% through 2031. This information shows the rising popularity of intelligent decision support systems in the field of healthcare. In general, RAG enhances the reasoning process by making it based on evidence, transparent, and context-sensitive.

Privacy, Accuracy, and Efficiency in Pediatric EHR Analysis

The results indicate that privacy and retention of clinically accurate longitudinal data are essential in successful analysis of pediatric EHRs. Pediatric EHRs include very sensitive data as children need consistent monitoring of their growth, immunization history, development status, and familial health data for several years. The RAG architecture increases privacy by collecting only necessary clinical data instead of accessing entire patient records during the analysis process (Sangaroonsilp *et al.*, 2021). The analysis has also revealed unique accuracy needs of pediatric EHRs compared to adult systems. First, growth evaluation is based on age-based percentile charts and z-scores in terms of height, weight, body mass index, and head circumference, instead of using constant references.




Figure 4: Pediatric EHR security

(Source: Itransition, 2023)

Moreover, the longitudinal EHR system must assess the growth rate between visits in order to determine any abnormal developments. Accuracy can also be increased through the use of immunization prediction, which verifies the previous vaccine administrations, identifies due or overdue immunizations and adjusts recommendations in cases of premature birth, immunodeficiency, or adverse reactions to vaccines. Developing EHR applications for pediatric populations allow for integration of information from developmental screening tests, pre-visit questionnaires, and social determinants of health into standardized reports that can be accessed later on. Standardized reporting allows managing of immunizations for populations, as it allows detection of unimmunized children and helps in following up with such patients (Atkinson *et al.*, 2020). Thus, privacy-preserving retrieval of data, pediatric-specific clinical computation, and longitudinal monitoring have been shown to increase the effectiveness and precision of pediatric EHR analysis.

Discussion

In this paper, the application of RAG architecture to analyze longitudinal EHRs in pediatric health care was considered. It was found that RAG provides an improvement in retrieving the relevant clinical information by combining historical

patient data with the evidence-based medical knowledge. Thus, RAG gives a better picture of the medical history of pediatric patients during different visits (Beede *et al.*, 2020). The results showed that continuity of care was supported in the context of integration of diagnoses, laboratory data, medication, immunization, and clinical notes into one timeline of patient information. The context-based retrieval allows the clinicians to interpret the disease progress, therapy response, and follow-up outcomes without searching for fragmented medical records. Such an approach is especially useful for children who have to be observed within long periods of time. Finally, it was revealed that RAG improved clinical decision support by providing responses based on verified clinical data (Swift *et al.*, 2021). Evidenced-based search enhances diagnostic reasoning and minimizes the chances of making recommendations that may not be sufficiently based on data. Structured search is efficient since it provides clinicians with the most pertinent data concerning patients' conditions. Nonetheless, there are barriers to the application of RAGs in pediatric healthcare facilities. For instance, fragmented data among various institutions, poor quality of documentation, and stringent privacy regulations may influence the results of search processes. Furthermore, pediatric EHRs require specific interpretation concerning development and use of medications among others (Ni *et al.*, 2021). On balance, RAG frameworks are a good tool for enhancing pediatric EHR analysis.

Conclusion

The paper proves that Retrieval-Augmented Generation models help analyze longitudinal pediatric EHRs in terms of retrieving appropriate medical records and linking medical files of patients across many health visits. The results indicate that RAG helps increase the context of patient history and clinical decision support. Moreover, the methodology makes it possible to be more effective since it eliminates the need to manually check fragmented medical records without compromising the privacy of patients. Taking into account the difficulties associated with integration of data and complexity of clinical issues regarding pediatric medicine, RAG is considered to be a promising solution for analyzing EHRs.

Reference List

1. Atkinson, K. M., Mithani, S. S., Bell, C., Rubens-Augustson, T., & Wilson, K. (2020). The digital immunization system of the future: imagining a patient-centric, interoperable immunization information system. *Therapeutic advances in vaccines and immunotherapy*, 8, 2515135520967203. <https://journals.sagepub.com/doi/abs/10.1177/2515135520967203>
2. Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., & Vardoulakis, L. M. (2020, April). A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In *Proceedings of the 2020 CHI conference on human factors in computing systems* (pp. 1-12). <https://dl.acm.org/doi/abs/10.1145/3313831.3376718>
3. Gumiel, Y. B., Silva e Oliveira, L. E., Claveau, V., Grabar, N., Paraiso, E. C., Moro, C., & Carvalho, D. R. (2021). Temporal relation extraction in clinical texts: a systematic review. *ACM Computing Surveys (CSUR)*, 54(7), 1-36. <https://dl.acm.org/doi/abs/10.1145/3462475>
4. Hashem, F., Brigden, C., Wilson, P., & Butler, C. (2020). Understanding what works, why and in what circumstances in hospice at home services for end-of-life care: applying a realist logic of analysis to a systematically searched literature review. *Palliative Medicine*, 34(1), 16-31. <https://journals.sagepub.com/doi/abs/10.1177/0269216319867424>
5. Hewa, T., Braeken, A., Ylianttila, M., & Liyanage, M. (2020, December). Multi-access edge computing and blockchain-based secure telehealth system connected with 5G and IoT. In *GLOBECOM 2020-2020 IEEE Global Communications Conference* (pp. 1-6). IEEE. <https://ieeexplore.ieee.org/abstract/document/9348125/>
6. Itransition, 2023. Pediatric EHR: key features, benefits, and leading platforms. <https://www.itransition.com/healthcare/ehr/pediatric>
7. Kovalchuk, S. V., Kopanitsa, G. D., Derevitskii, I. V., Matveev, G. A., & Savitskaya, D. A. (2022). Three-stage intelligent support of clinical decision making for higher trust, validity, and explainability. *Journal of Biomedical Informatics*, 127, 104013. <https://doi.org/10.1016/j.jbi.2022.104013>
8. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33, 9459-9474. https://proceedings.neurips.cc/paper_files/paper/2020/hash/6b493230-Abstract.html

9. Ni, Y., Bachtel, A., Nause, K., & Beal, S. (2021). Automated detection of substance use information from electronic health records for a pediatric population. *Journal of the American Medical Informatics Association*, 28(10), 2116-2127. <https://academic.oup.com/jamia/article-abstract/28/10/2116/6333356>
10. Oelsner, E. C., Balte, P. P., Bhatt, S. P., Cassano, P. A., Couper, D., Folsom, A. R., ... & Yende, S. (2020). Lung function decline in former smokers and low-intensity current smokers: a secondary data analysis of the NHLBI Pooled Cohorts Study. *The Lancet Respiratory Medicine*, 8(1), 34-44. [https://www.thelancet.com/journals/lanres/article/PIIS2213-2600\(19\)30276-0/abstract](https://www.thelancet.com/journals/lanres/article/PIIS2213-2600(19)30276-0/abstract)
11. Pai, M. M., Ganiga, R., Pai, R. M., & Sinha, R. K. (2021). Standard electronic health record (EHR) framework for Indian healthcare system. *Health Services and Outcomes Research Methodology*, 21(3), 339-362. <https://link.springer.com/article/10.1007/s10742-020-00238-0>
12. Pureinsights, 2023. What is Retrieval Augmented Generation (RAG)?
13. Rosset, C., Xiong, C., Phan, M., Song, X., Bennett, P., & Tiwary, S. (2020). Knowledge-aware language model pretraining. *arXiv preprint arXiv:2007.00655*. <https://arxiv.org/abs/2007.00655>
14. Sangaroonsilp, P., Dam, H. K., Choetkiertikul, M., Ragkhitwetsagul, C., & Ghose, A. (2021). Mining and classifying privacy and data protection requirements in issue reports. *arXiv preprint arXiv:2112.13994*. <https://arxiv.org/abs/2112.13994>
15. Shang, Y., Tian, Y., Zhou, M., Zhou, T., Lyu, K., Wang, Z., ... & Li, J. (2021). EHR-oriented knowledge graph system: toward efficient utilization of non-used information buried in routine clinical practice. *IEEE Journal of Biomedical and Health Informatics*, 25(7), 2463-2475. <https://ieeexplore.ieee.org/abstract/document/9444689/>
16. Sinhababu, N., Saxena, R., Sarma, M., & Samanta, D. (2021). Medical information retrieval and interpretation: a question-answer based interaction Model. *arXiv preprint arXiv:2101.09662*. <https://arxiv.org/abs/2101.09662>
17. Swift, J., Harris, Z., Woodward, A., O'Kelly, N. I., Barker, C., & Ghosh, S. (2021). The implementation of a virtual ward using digital solutions informing community clinicians in early supported discharge of patients with SARS-Cov2 respiratory symptoms from an acute hospital setting. *medRxiv*, 2021-03. <https://www.medrxiv.org/content/10.1101/2021.03.29.21254548.abstract>
18. Vest, J. R., Unruh, M. A., Casalino, L. P., & Shapiro, J. S. (2020). The complementary nature of query-based and directed health information exchange in primary care practice. *Journal of the American Medical Informatics Association*, 27(1), 73-80. <https://academic.oup.com/jamia/article-abstract/27/1/73/5583717>
19. Zhang, D., Yin, C., Zeng, J., Yuan, X., & Zhang, P. (2020). Combining structured and unstructured data for predictive models: a deep learning approach. *BMC medical informatics and decision making*, 20(1), 280. <https://link.springer.com/article/10.1186/s12911-020-01297-6>
20. Zhu, F., Lei, W., Wang, C., Zheng, J., Poria, S., & Chua, T. S. (2021). Retrieving and reading: A comprehensive survey on open-domain question answering. *arXiv preprint arXiv:2101.00774*. <https://arxiv.org/abs/2101.00774>