International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 10 | Oct - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Apha: AI-Powered Personalized Healthcare Assistant — A Review of Personal Health Agents, Machine-Learning Risk Prediction, Medical Report OCR, and Retrieval Augmented Biomedical Question Answering

Under the Guidance of **Prof Aparna R.Khairkar , Prof Monika S. Shirbhate**

Anvi Mankar (38), Prachi Uttarwar (77),Diya Dubey (71), Suhani Raut (46)
Department of Information Technology, Prof. Ram Meghe College of Engineering and Management, Badnera, Amravati, Maharashtra, India | Academic Session 2026–27

I ABSTRACT

Taking care of your own health is a daily job. You have to keep your reports in order, understand what the numbers mean, remember to take medicines on time, keep an eye on long-term conditions, and decide which health advice is worth trusting. In the last few years, progress in artificial intelligence (AI), machine learning (ML), large language models (LLMs), optical character recognition (OCR) and retrieval-augmented generation (RAG) has made it possible to help with many of these tasks. In this paper we review twelve studies that cover four areas: personal health agents and LLM-based coaching, ML models for disease risk prediction, OCR for medical reports, and RAG for biomedical question answering. Our main observation is that each area works well on its own, but they are rarely put together. Existing systems also tend to explain things poorly to ordinary patients, ignore the user's own records, and say little about safety, fairness and keeping users interested over time. Based on this, we describe a concept called the AI Personalized Healthcare Assistant (APHA). It is one app that offers chat support, report analysis, disease risk estimates, personal lifestyle advice, medicine reminders and secure health records. We see APHA as a tool that helps doctors and patients, not one that replaces doctors. The biggest challenges we found are data privacy, bias, wrong answers from LLMs, and user engagement.

Keywords: personalized healthcare, artificial intelligence, large language models, disease risk prediction, medical report analysis, retrieval-augmented generation, OCR.

II. INTRODUCTION

Looking after your health is harder than it sounds. A person has to keep medical records, read test reports, remember medicines, track conditions like diabetes or blood pressure, and figure out which information online is reliable. Most people end up using one app for fitness, another for reminders, a third for reports, and a website for health information. The result is that their health data is spread out and hard to follow over time. This is even harder for older people, for patients with more than one condition, and for anyone who finds medical language confusing.

AI has moved quickly in the meantime. Machine learning, computer vision and LLMs can now handle both structured data (like lab values) and unstructured data (like free text and images), and they can answer medical questions in plain language. LLMs in medicine have received a lot of attention [1], and broader reviews describe both big opportunities and real risks in biomedicine and health [2]. Reviews of LLMs in personalized healthcare [9] and of LLM agents in biomedicine [12] show that the field is moving toward personal and agentbased support, but they also stress that these systems need careful validation and human oversight. The wider role of AI in medicine has been discussed in the medical community as well [11]. Work on medical question answering also shows that LLMs hold a fair amount of clinical knowledge, but that careful human evaluation still finds important gaps [13].

The trouble is that most of these systems work alone. Some give wellness coaching [3], [4], some predict disease risk [7], some digitize reports [6], and some answer questions from medical papers [5]. Only a few try to do several of these things together. A truly useful health assistant should do more than produce text. It should read the user's own data, explain it in simple words, estimate risk carefully, suggest practical steps, remind the user about medicines, protect privacy, and know when to say "please see a doctor".

This paper reviews the main research behind such assistants and then uses APHA as an example of a single, patient-centred platform. We are not only summarizing studies. We also want to show how the strengths and weaknesses of each study shape the design of a combined system.

A. Objectives of the Review

- To look at studies on personal health coaching and conversational health guidance.
- To see how ML models are used for disease risk prediction and how their results can be made understandable for patients.
- To study how OCR and layout analysis turn scanned or photographed reports into structured data.

- To check how RAG can make biomedical answers better supported by evidence.
- To find the gaps in current work that lead to the APHA concept.

III. REVIEW METHODOLOGY

This is a technology-focused narrative review, organized around the parts of the proposed APHA system. We first collected recent and relevant papers from journals, conferences and preprint servers. We looked mainly at peer-reviewed work and research results on LLMs in medicine, personal health agents, ML-based risk prediction, OCR for medical reports, and retrievalaugmented biomedical question answering. The final list has twelve core works [1]–[12], published mostly between 2023 and 2025. We also use five supporting references [13]– [17] that cover model evaluation, hallucination, algorithmic bias and health-data standards.

We sorted the papers into five groups: (1) personal health agents and LLM coaching, (2) ML for disease risk prediction, (3) OCR and digitization of medical reports, (4) retrieval-augmented biomedical question answering, and (5) general surveys and opinion pieces on LLMs and AI in medicine. The grouping follows each paper's main contribution, how useful it is for APHA, the limits the authors mention, and the points that matter when building a complete patient-facing system.

We did not try to list every AI healthcare technology that exists. Our aim was simpler: to collect enough evidence to show that the APHA idea is technically realistic, and to understand what is needed for its implementation, safety and testing.

IV. LITERATURE REVIEW

A. Personal Health Agents and LLM-Based Coaching Some recent work builds personal health systems from several LLMbased agents, each with its own job, such as analysing data, giving medical knowledge, or acting as a coach [4]. These agents use data from wearables and biomarkers to give personal feedback and lifestyle tips. This shows that LLMs can manage several analysis tools and keep up a long conversation. But they mostly stay in the wellness area. They do not really handle medical report analysis, clear disease risk prediction or medication management.

Other researchers built personal health LLMs tuned for specific areas such as sleep and fitness, and on benchmark tasks these models often reach expert-level performance [3]. Fine-tuning a model for one domain makes its advice better and safer. The downside is that these models remain narrow. They rarely deal with things like reading lab values or looking after a patient who has several conditions at once.

Frameworks for conversational health agents show how different kinds of data and tools can be connected inside multistep conversations [8]. Even so, reading reports with OCR and giving users clear risk levels is still not well covered. For APHA, the lesson is that a modular design, where the chatbot coordinates specialized tools, is a practical way to build a broad assistant.

B. Machine Learning for Disease Risk Prediction

ML models are widely used to predict the risk of heart disease, diabetes and kidney disease. One example is the use of gradient boosting methods such as LightGBM on routine clinical data to estimate the chance of major adverse cardiac events in high-risk groups, with good results [7]. Models like

this help sort patients by risk and decide how closely each person should be followed up.

The catch is that these predictors mostly stay in research projects or in tools made for doctors. They are rarely combined with a chat interface, lifestyle coaching and record keeping in one app for patients. For APHA, this means that a trained model is only one piece. Its output also has to be explained, shown with its uncertainty, and tied to sensible next steps.

C. OCR and Digitization of Medical Reports

A lot of important lab and imaging results are still stored as scanned or photographed reports. End-to-end OCR and layoutanalysis pipelines can rebuild these report images and turn them into tables a computer can read, which opens the door to further analysis and decision support [6].

Most of this work, though, stops at digitization. A patient still needs help to know whether the values are normal, what a change over time means, and when to see a doctor. Filling this gap needs OCR to be combined with explanation, risk estimation and long-term tracking. The extracted values must also be checked, because one wrongly read digit in a lab result can change its meaning completely.

D. Retrieval-Augmented Biomedical Question AnsweringRAG systems combine information retrieval with LLMs so that answers rest on current medical literature. Pipelines built at PubMed scale, using BM25 retrieval and specialized rerankers, give accurate answers with practical response times [5]. Surveys on LLMs in biomedicine see strong potential here for grounded question answering, summarization and guideline checking [2], and a systematic review of RAG in healthcare comes to a similar view [10]. A medical RAG benchmark called MIRAGE, with 7,663 questions from five datasets, was built to test how different corpora, retrievers and LLMs behave together, and it reports that RAG can improve the accuracy of the LLMs tested [16].

These methods are very good at giving evidence-based answers, but they usually work with general literature. They do not bring in a person's own records, reports and risks. So in APHA the retrieval layer should search trusted medical sources, while the final answer should also take the user's own data into account and keep the source of each statement visible.

D. Surveys and Perspectives on LLMs in Medicine General surveys describe LLMs as useful for answering medical questions, summarizing information and talking to patients [1], [2]. Reviews that focus on personalization and agents [9], [12] discuss how models can be adapted to individual patients and how autonomous agents can use tools. The perspective article on AI in medicine [11] reminds readers that clinical value has to be shown, not assumed. In the same spirit, the evaluation of MedPaLM and Flan-PaLM [13] proposes judging model answers on several axes, including factuality, possible harm and bias, and not only on benchmark accuracy. Taken together, these works give the basic principles for APHA. They also warn about madeup answers, biased outputs, privacy risks and the need for human oversight.

VI. RESEARCH GAPS IDENTIFIED

The literature points to several gaps, and these are what motivate the APHA concept.

1) Fragmented Functionality = Coaching agents, risk predictors, OCR pipelines and RAG systems are usually built and used separately. A user has to jump between tools and carry information from one to another by hand. Very few platforms bring chatbot support, OCR-based report analysis, diseasespecific risk prediction, personal advice, medicine reminders and secure record storage into one easy-to-use assistant. APHA is meant to reduce this effort.

COMPARATIVE ANALYSIS solve a single sub-problem well, but putting everything into one Table I shows how each surveyed study relates to the APHA safe, patient-friendly system is still a major engineering and workflow. The comparison makes one thing clear: most studies research challenge.

TABLE I. COMPARISON OF SURVEYED STUDIES AND THEIR RELEVANCE TO APHA

Study	Area	Core contribution	APHA relevance	Main limitation
Thirunavukarasu et al. [1]	LLMs in medicine	Reviews how LLMs may be used in medical tasks and what risks they bring	Gives the broad context for using LLMs in a health assistant	General overview; not a patient-facing system design
Tian et al. [2]	LLMs in biomedicine	Discusses opportunities and challenges of ChatGPT-type models in biomedicine and health	Highlights benefits and risks to consider for APHA	Broad scope; no integrated workflow
Cosentino et al. [3]	Personal health LLM	Tunes an LLM for sleep and fitness coaching with expert-level benchmark results	Supports personalized lifestyle advice	Narrow focus on sleep and fitness
Xu & Heydari [4]	Personal health agent	Describes a multi-agent design with data analysis, medical expertise and coaching roles	Supports modular agents for APHA	Stays mostly within wellness; limited report and risk analysis
Stuhlmann et al. [5]	Biomedical RAG	Builds PubMed-scale RAG with BM25 retrieval and rerankers	Provides evidence-based answers for APHA	Uses general literature, not the user's own records
Ren et al. [6]	Report OCR	Digitizes image-based lab reports into machine-readable form	Enables report reading in APHA	Stops at digitization; no explanation or risk
Kaczmarek et al. [7]	Risk prediction	Uses ML to predict major adverse cardiac events in diabetes and kidney disease	Supports disease risk estimation	Built for clinical settings, not for patients
Abbasian et al. [8]	Conversational health agents	Proposes an LLM-powered framework combining data, tools and dialogue	Provides a model for the APHA chat engine	Limited OCR and clear risk levels
Souza et al. [9]	LLMs in personalized care	Systematically reviews LLMs for tailoring treatment and care	Guides personalization in APHA	Review only; no implemented platform
RAG in healthcare review [10]	RAG review	Reviews retrieval-augmented LLMs in healthcare	Guides grounded answering	Does not cover personal-data grounding
Beam et al. [11]	AI in medicine	Presents a perspective on the role of AI in medicine	Reminds that clinical value must be shown	Perspective rather than technical method
Xu & Sankar [12]	LLM agents review	Reviews LLM agents for biomedicine	Supports agent-based design	Broad; few patient-facing evaluations

2) Digitization Without Interpretation = OCR and layoutanalysis work can turn a report into a table, but the patient still does not know if the values are normal, how they have changed, or when a doctor is needed. Interpretation, comparison with reference ranges and trend tracking must be added on top of digitization.

3) Limited Patient-Facing Risk Communication = Risk predictors are mostly designed for research or for doctors. Patients need risk scores with plain-language explanations, the main contributing factors, a note about uncertainty, and a clear reminder that the score is not a diagnosis.

4) Weak Personal Grounding and Traceability = RAG systems give evidence-based answers from general literature, but they usually do not link those answers to a person's own reports and risks. It should also be easy to trace an important statement back to its source, because an AI answer can sound convincing even when the supporting evidence is incomplete.

5) Safety, Privacy and Fairness = LLMs can give answers that sound right but are wrong, a problem usually called hallucination [14]. Health data is highly sensitive, and models trained on unrepresentative data may work badly for some groups. Safety layers, strong security and bias checks are therefore needed.

6) Engagement and Clinical Integration = Even a welldesigned assistant is of no use if people stop using it. And an assistant that cannot connect with electronic health records or telemedicine platforms will not fit into real clinical workflows.

7) Lack of Standardized Evaluation = There is no single measure that tells us whether a healthcare assistant works properly. OCR accuracy, risk-model quality, answer faithfulness, safety behaviour and user adherence all need to be checked separately and then together.

VII. PROPOSED APHA FRAMEWORK

APHA is meant to be one platform for everyday health management. Its main goals are to:

- Give conversational, AI-powered health guidance suited to each user.
- Create personal diet and exercise plans based on a person's health status and goals.
- Analyse uploaded medical reports using OCR and AI-based interpretation.
- Estimate disease risks with ML models and point out the key risk factors.
- Help users stay regular with medicines through reminders and alerts.
- Show important indicators and trends on an easy-to-read health dashboard.

The assistant is meant to support qualified healthcare professionals, not to replace them.

A. System Architecture

A simplified APHA architecture has these components:

- **User Interface:** A web or mobile app where users chat with the assistant, upload reports, enter data and view dashboards.
- **Conversational Engine:** An LLM-based chatbot that understands questions, collects structured information and coordinates the other modules.
- **OCR and Report Analysis:** A pipeline using tools like Tesseract or EasyOCR, plus NLP, to turn images or PDFs into structured values and short explanations.
- **Risk Prediction Models:** ML models trained on relevant clinical features that estimate risk for selected diseases and explain what contributes to it.
- **Lifestyle Recommendation Engine:** Modules that build diet and exercise plans using age, existing conditions and personal preferences.
- **RAG-Based Medical QA:** A retrieval layer over trusted sources, combined with an LLM, to give well-grounded answers.
- **Reminder and Alert System:** Scheduling services that send medicine reminders and time-sensitive alerts.
- **Secure Data Storage:** Encrypted databases and file storage with strong authentication and audit logs.

B. Conceptual Workflow

User Login and Consent → Chat Query or Report Upload → OCR and Layout Analysis → Value Extraction and Normalization → Comparison with Reference Ranges → Disease Risk Estimation → Evidence Retrieval (RAG) → Personalized Advice Generation → Safety Check → Dashboard, Reminders and Alerts.

C. Report Interpretation and Risk Estimation Strategy The report module first pulls out each test name, value, unit and reference range from the uploaded report and checks them for obvious errors. Next, it compares every value with its reference range and marks it as normal, low or high. It then compares the current values with the user's earlier reports to show trends. After that, the verified values, along with details the user entered (such as age, history and lifestyle), go to the relevant diseasespecific ML model. Finally, the result is shown with the main contributing factors, how uncertain the estimate is, and a suggestion to see a doctor when needed. The user should be able to correct any value that was read wrongly.

D. Safety-Aware and Evidence-Grounded Generation Answers should be based on trusted medical sources fetched by the RAG layer and on the user's own verified data. Simple rulebased checks should run before an answer is shown, for example to catch emergency symptoms, unsafe medicine advice, or claims that go beyond the available evidence. When there is not enough supporting information, the system should

say so and advise the user to talk to a healthcare professional rather than guess.

VIII. KEY FUNCTIONALITIES

A. AI Chatbot and Symptom Guidance

The chatbot is the main entry point. It asks about symptoms, history and lifestyle, explains medical terms, and guides users to features like report upload or risk assessment. To avoid harm, it should give careful, conservative advice and clearly tell users to see a doctor in emergencies or when symptoms are unclear.

B. Medical Report Interpretation

Using OCR and NLP, the assistant reads values from lab reports, compares them with reference ranges and summarizes the results. It can flag high glucose or cholesterol and explain in simple words what that means. This helps patients move from raw numbers to something they can act on.

C. Disease Risk Estimation

By combining details entered by the user with values taken from lab reports, the system can run disease-specific ML models to estimate risk, such as the chance of a heart-related event. Each risk score should come with a clear explanation, information about uncertainty, and a reminder that it is a supporting tool and not a diagnosis.

D. Lifestyle and Medication Support

The assistant creates individual diet and exercise plans based on the user's goals (for example, weight control or managing blood pressure), health conditions and cultural preferences. A reminder module sends notifications for medicines and appointments and tracks how regularly the user follows the routine, which they can view over time.

E. Dashboards and Additional Services

A dashboard brings together vital signs, lab trends and adherence data. Optional extras include a doctor-finder, moderated community groups for peer support, and quick access to emergency contacts and local services.

IX. TECHNOLOGY STACK AND IMPLEMENTATION CONSIDERATIONS

The choices below are only suggestions and would suit a prototype of a healthcare assistant:

- **Frontend:** A web or mobile interface with a chat window, report upload, dashboard charts, reminders and a settings page for privacy and consent.
- **Backend:** Login and authentication, API handling, OCR and analysis services, risk-model serving, retrieval services, reminder scheduling and audit logging.
- **AI and NLP:** OCR with Tesseract or EasyOCR, NLP for value extraction, an LLM for conversation and explanation, and gradient-boosting or similar ML models for risk estimation.
- **Retrieval:** A trusted medical knowledge base with keyword (BM25) and semantic search, optionally with a reranker, as in [5].
- **Storage and Standards:** Encrypted databases and file storage, with standards such as FHIR [17] for connecting to other health systems.
- **Evaluation Data:** A carefully chosen set of sample reports, questions and cases, labelled by clinicians, to test each component.

X. EVALUATION FRAMEWORK

To judge APHA properly, its results should be compared with baselines and with expert opinion. OCR and extraction quality can be measured by field-level accuracy, precision, recall and F1 score for items such as test name, result, unit and reference range. Risk models can be assessed with measures like AUROC, calibration, and sensitivity and specificity on suitable datasets. The answers from the RAG component can be judged for correctness, faithfulness to the retrieved sources and relevance of citations. Safety can be tested with a set of risky prompts, such as emergency symptoms or requests for unsafe medicine advice, to see whether the assistant responds conservatively.

Lifestyle plans and explanations should be reviewed by qualified healthcare professionals for safety, clarity and suitability. On the user side, we should measure usability and satisfaction scores, how well people follow medicine reminders, and whether they keep using the app over time, since APHA is meant for regular use.

Finally, an ablation study should show what each component contributes. For example, chatbot answers can be compared with and without RAG, risk explanations with and without uncertainty information, and one general chatbot against the modular design.

XI. CHALLENGES, ETHICS AND LIMITATIONS

Using AI in personal healthcare brings risks that should be dealt with openly.

Privacy and security. Health data must be encrypted, accesscontrolled and handled according to the relevant regulations. Users should give informed consent and be able to control and delete their data. One breach can destroy trust very quickly.

Reliability of LLMs. Text generation models are known to produce hallucinated text, meaning content that is not supported by the source [14], and LLMs can therefore give answers that sound right but are wrong when they are not kept in check. Safety layers, rule-based checks and grounding in trusted sources are therefore essential.

Bias and fairness. Models trained on unrepresentative data may work poorly for certain groups and make existing inequalities worse. A well-known example is a widely used health-risk algorithm that showed racial bias because it used healthcare costs as a stand-in for health needs [15]. Bias audits, diverse datasets and continuous monitoring are important.

User engagement. Even a well-designed assistant is useless if people stop using it. Ideas from behavioural science and usability research can guide features such as reminders, goalsetting and feedback that keep users coming back.

Clinical integration. The assistant has to work with electronic health records and telemedicine platforms if it is to fit into real clinical workflows. Following standards such as FHIR, which was created to make healthcare information exchange easier [17], and designing together with clinicians can help APHA support medical practice rather than disturb it.

Responsibility and limits. Risk scores and advice must never be presented as a diagnosis or as a replacement for a doctor. One limitation of this paper is that the framework has not been tested as a complete system. Our claims come from a design based on the literature, not from experimental results. When the system is built, the datasets, model versions, evaluation methods, error analysis and repeatable experiments should be properly documented.

XII. FUTURE SCOPE

APHA can be developed further in several ways. It could use data from wearable devices, support more languages and voice interaction, and understand medical images as well as text reports. Personalization could improve if the system remembers a user's history over a longer period, and privacy-preserving learning methods could allow models to be trained without gathering sensitive data in one place. Closer links with electronic health records and telemedicine platforms, along with a view for clinicians, would let doctors review what the assistant has collected. Explainable risk estimation is another useful direction: users should be able to see which factors raised or lowered their risk, how certain the estimate is, and which sources support the advice.

XIII. CONCLUSION

AI-driven personalized healthcare assistants are a promising way to manage everyday health. Existing research already offers strong building blocks: personal health agents, ML-based risk prediction, OCR-based report digitization and RAG-based biomedical question answering. However, most current systems are narrow and disconnected, and important problems remain in integration, patient-friendly interpretation, use of personal data, safety, fairness and evaluation.

The APHA concept shows how these pieces can be joined into one patient-centred platform that offers conversational guidance, report analysis, risk estimation, lifestyle recommendations, medicine reminders and secure record storage. In the future, such systems must be private, clinically reliable, fair, engaging and closely connected with healthcare professionals, so that they strengthen human care instead of replacing it. The next step is to build a working prototype and test it against clear benchmarks.

REFERENCES

[1] A. J. Thirunavukarasu, D. S. J. Ting, K. Elangovan, L. Gutierrez, T. F. Tan, and D. S. W. Ting, "Large language models in medicine," Nature Medicine, vol. 29, pp. 1930–1940, 2023.

[2] S. Tian, Q. Jin, L. Yeganova, et al., "Opportunities and challenges for ChatGPT and large language models in biomedicine and health," Briefings in Bioinformatics, vol. 25, bbad493, 2024.

[3] J. Cosentino, et al., "A personal health large language model for sleep and fitness coaching," Nature Medicine, vol. 31, pp. 3394–3403, 2025.

[4] X. Xu and A. Heydari, "The anatomy of a personal health agent," Google Research, 2025.

[5] L. Stuhlmann, M. A. Saxer, and J. Fürst, "Efficient and reproducible biomedical question answering using retrievalaugmented generation," arXiv preprint, 2025.

[6] X. Ren, D. Dou, X. He, F. Cui, and J. Zhao, "A serialization method for digitizing the image-based medical laboratory report," Digital Health, 2025.

[7] M. Kaczmarek, et al., "Predicting major adverse cardiac events in diabetes and chronic kidney disease," Cardiovascular Diabetology, vol. 24, art. 76, 2025.

[8] M. Abbasian, I. Azimi, A. M. Rahmani, and R. Jain, "Conversational health agents: A personalized LLM-powered agent framework," arXiv preprint, 2023.

[9] G. d. P. Souza, G. Melo, and D. Schneider, "Tailoring treatment in the age of AI: A systematic review of large language models in personalized healthcare," Informatics, vol. 12, art. 113, 2025. [10] "Retrieval-augmented generation for large language models in healthcare: A systematic review," PLOS Digital Health, 2025.

[11] A. L. Beam, J. M. Drazen, I. S. Kohane, et al., "Artificial intelligence in medicine," New England Journal of Medicine, vol. 388, pp. 1220–1221, 2023.

[12] X. Xu and R. Sankar, "Large language model agents for biomedicine: A comprehensive review," Information, vol. 16, art. 894, 2025.

[13] K. Singhal, S. Azizi, T. Tu, et al., "Large language models encode clinical knowledge," Nature, vol. 620, no. 7972, pp. 172–180, 2023, doi: 10.1038/s41586-023-06291-2.

[14] Z. Ji, N. Lee, R. Frieske, T. Yu, D. Su, Y. Xu, E. Ishii, Y. Bang, A. Madotto, and P. Fung, "Survey of hallucination in natural language generation," ACM Computing Surveys, vol. 55, no. 12, art. 248, pp. 1–38, 2023.

[15] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, "Dissecting racial bias in an algorithm used to manage the health of populations," Science, vol. 366, no. 6464, pp. 447–453, 2019,doi: 10.1126/science.aax2342.

[16] G. Xiong, Q. Jin, Z. Lu, and A. Zhang, "Benchmarking retrievalaugmented generation for medicine," in Findings of the Association for Computational Linguistics: ACL 2024, Bangkok, Thailand, 2024, pp. 6233–6251, doi: 10.18653/v1/2024.findingsacl.372.

[17] D. Bender and K. Sartipi, "HL7 FHIR: An agile and RESTful approach to healthcare information exchange," in Proc. 26th IEEE Int. Symp. Computer-Based Medical Systems (CBMS), 2013, pp. 326– 331, doi: 10.1109/CBMS.2013.6627810.

Prepared as an academic review-paper draft for the APHA project.