

A Review Paper on an Integrated Multi-Agent Decision Support System for Medical Image Analysis, Clinical Diagnosis, Drug Interaction Assessment, and Risk Evaluation

Vaishnavi Tushar Jagade, Saniya Salim Bijali, Diya Ramesh Dawle, Anushka Kiran Bhosale, Priti Vanmali

Department of Computer Engineering, Navsahyadri Education Society's Group of Institutions, Pune, Maharashtra, India

Published October 2026

Abstract

Healthcare decision making often requires multiple doctors to consider different types of information like patient symptoms, their medication history, medical images, lab reports, and relevant medical literature. All this information is frequently fragmented across multiple sources. A multi-agent decision support system can help to integrate all these sources and organize them. This paper presents the design and development of a Multi-Agent Healthcare Decision Support System (DSS) designed to integrate medical image analysis, clinical diagnosis, drug-interaction assessment, evidence retrieval, and risk evaluation through a modular architecture.

The proposed system consists of 6 different specialized agents namely vision agent, diagnosis agent, drug-interaction agent, literature/RAG agent, emergency risk agent and Coordinator/Orchestrator. A single agent is not able to do all tasks. Image analysis, medication assessment, evidence retrieval, and risk evaluation have different computational requirements and need different validation procedures for each agent.

The system's effectiveness must be established through systematically designing, developing and testing the individual agent and integrating the whole workflow. The proposed system aims to support clinical decision making, not to replace professionals.

Keywords: multi-agent system; clinical decision support system; medical image analysis; drug interaction; retrieval-augmented generation; risk assessment

1 Introduction

Healthcare professionals continuously need to evaluate different types of medical information before making any decision. These may include patient symptoms, previous diagnoses, medical images, medication reports, and laboratory results. Reviewing all these sources independently can make the information integration hard.

Artificial intelligence (AI) offers different approaches to analyse all these sources and organise decision relevant information. Different systems use AI only for one agent specification, but in the multi-agent DSS AI helps according to each agent individually. A single model is not sufficient for executing all the tasks, so the different models do these tasks very efficiently.

In such systems, individual agents are assigned the different responsibilities and communicate through the coordinating components. This architecture can improve modularity, support independent testing and make it easier to extend, validate, and maintain individual components without

redesigning the entire decision pipeline. This system is designed to provide well organized and evidence-based outputs for review by multiple healthcare professionals.

This system saves the time of healthcare specialists by giving integrated information to clinicians in a structured, traceable format, so that each recommendation can be reviewed alongside the underlying imaging findings, medication profile, and supporting evidence.

1.1 Problem statement

Healthcare information is distributed across different formats and sources. Medical image analysis, medication safety assessment, biomedical literature retrieval, and risk evaluation may require separate tools or workflows. When these functions are not coordinated, users may need to switch between systems and manually combine their outputs.

The research problem is to design a modular AI-based decision support framework that coordinates specialized analytical components and presents their outputs in a structured manner while maintaining appropriate human oversight.

1.2 Objectives

The principal objectives are:

1. To design a modular multi-agent architecture for healthcare decision support.
2. To develop a vision agent for medical image analysis.
3. To develop a diagnosis support agent that processes relevant patient information.
4. To implement drug interaction analysis using structured drug information.
5. To develop a retrieval-augmented generation component for biomedical literature retrieval.
6. To design a risk assessment component for identifying potentially concerning clinical information.
7. To coordinate the outputs of specialized agents through a central orchestration mechanism.
8. To evaluate the system using task-specific performance, reliability, and safety metrics.
9. To provide an interface that presents findings and supporting information for human review.

1.3 Organisation of the paper

The remainder of the paper reviews related work and identifies the problems addressed by the project, describes the proposed system architecture, explains the methodology and the individual agents, and describes the implementation of the DSS. It then presents the current status and evaluation plan, discusses the proposed system and its limitations, and concludes with directions for future work.

2 Related work

2.1 AI in healthcare decision support

AI-based decision-support systems can assist with relevant information analysis, clinical workflow support, and the presentation of different findings. The usefulness depends on the input data which we are using, its quality, the reliability of the models and the appropriateness of the recommendations.

The proposed system uses a modular architecture, with specialised components handling distinct healthcare tasks rather than relying on a single general-purpose model.

2.2 Drug information and medication safety

Drug information systems provide information about medicines, including warnings, different drug-label details, generic names, product information etc. OpenFDA provides programmatic access to selected U.S. Food and Drug Administration datasets, including drug-label information.

In the proposed system, the drug analysis module uses available drug data for medicine identification and the drug retrieval. Medicine-name normalisation is also important because a medicine may be referred to by a generic name, brand name, salt form, or regional name. Reliable identification requires more than simple text matching.

2.3 Retrieval-augmented generation and medical literature

The proposed system combines the retrieval-augmented generation (RAG) approach with a curated biomedical literature index, so that generated explanations are grounded in retrievable sources rather than in the language model's unverified parametric knowledge.

For this project the RAG workflow uses different medical literature and health-related documents collected through PubMed and WHO resources. The document-processing pipeline prepares retrieved material for searching and returns relevant passages for healthcare-related queries.

The retrieved literature should be presented as supporting evidence, not as proof that a particular treatment is appropriate for an individual patient.

2.4 Medical image analysis

The medical image analysis uses Machine Learning techniques and the computer vision to identify the patterns in the medical images. X-ray analysis may find the relevant outcomes like fractures, chest abnormalities and other related information depending on the data provided and the trained model.

2.5 Research gap

Healthcare related tools provide different individual tasks like, medicine information, medical image interpretation or different literature search. Integrating all these modules into one system is challenging. Communication between modules, consistency of results, evidence traceability, and safe presentation of findings are very challenging to manage efficiently.

This project addresses the integration challenge by providing a common backend and a framework which provides the coordination of specialized healthcare modules. The effectiveness of this project must be satisfied through the testing rather than the assumptions of the system design alone.

Table 1 summarises the reviewed work and the research gaps.

Table 1. Summary of reviewed healthcare decision support systems and research gaps.

Work	Focus	Limitation for This Problem
Sutton et al. [1]	Overview of clinical decision support systems	Limited discussion of integrating multiple AI healthcare modules
Rajpurkar et al. [2]	Deep learning for medical image analysis	Focuses on medical imaging rather than an integrated healthcare workflow

Work	Focus	Limitation for This Problem
Vamathevan et al. [3]	Machine learning applications in drug discovery and development	Does not provide a unified patient-level decision support system
Lewis et al. [4]	Retrieval-Augmented Generation (RAG)	General RAG framework; not specifically a multi-agent healthcare DSS
Jiang et al. [5]	Artificial intelligence in healthcare	Broad review; does not implement the complete proposed workflow
Esteva et al. [6]	Deep learning in healthcare	Focuses on AI applications rather than integrating all the proposed modules

3 Proposed system architecture

3.1 Overall structure

The Multi-Agent AI-Based Healthcare DSS is organised around a common application interface and backend API. User inputs are passed to the relevant healthcare modules, and their outputs are organized for presentation.

The overall architecture is shown in Figure 1.




Figure 1. System architecture of the multi-agent AI-based healthcare decision support system (DSS).

The architecture includes the following major components:

- *User Interface*: Collects patient information, medicine queries, symptoms, laboratory information, and medical image uploads.
- *Backend API*: Receives requests and connects the application to the relevant processing modules.
- *Symptoms and Laboratory Module*: Processes available symptom and laboratory information to support the proposed assessment workflow.
- *Vision Agent*: Processes supported X-ray images using the implemented image-analysis model.
- *Drug Information Agent*: Retrieves medicine-related information from available drug data sources.
- *RAG and Evidence Retrieval Module*: Searches indexed medical documents and returns relevant evidence.
- *Risk Assessment Module*: Evaluates supported risk indicators using implemented rules or analysis methods.
- *Results Interface*: Presents available module outputs and supporting information in an organised form.

The final implementation should distinguish between modules that are operational, partially implemented, and still proposed.

3.2 Processing workflow

Figure 2 illustrates the end-to-end processing workflow for a given user request, highlighting the parallel module invocations and the final user review.




Figure 2. Processing workflow for healthcare requests, from user input and validation to parallel specialised module analysis and final user review.

3.3 Data flow and storage

The system uses backend API to exchange information between applications and its processing module. In this system the Drug-interaction agent and the RAG agent use the respective data sources and the prepared documents.

Individual patient information and the uploaded reports should be handled according to the privacy controls.

4 Methodology

4.1 Input processing

The system accepts the medical input from the user through the application interface and the API. Depending on the user the input may include patient details, symptoms, laboratory reports, medicine information, X-ray images or medical related questions.

The backend should confirm the required fields and supported input formats before passing a request to the corresponding module. Invalid or unsupported inputs should produce an appropriate error response.

4.2 Drug information agent

The Drug Information Agent is designed to identify different medicines and to retrieve their associated information from available drug data sources. The agent is intended to support queries on medicine names, dosage forms, indications, contraindications, and known interactions, returning structured results that the Drug-Interaction assessment can consume downstream.

The current project implementation includes the drug analysis and uses OpenFDA for retrieving available drug-label information.

4.3 Retrieval-augmented generation module

The RAG component supports different medical documents and uses a relevant passage to support the healthcare related queries. The current project includes the documents related to the healthcare associated with PubMed and WHO.

The retrieval quality depends on the document coverage, text-extraction quality, query formulation and the retrieval methods.

4.4 Vision agent

The vision agent is intended to process and produce image-analysis outputs. This agent also gives output for X-ray images showing the related appropriate disease. It extracts image related information and passes it to the diagnosis agent and other components of DSS. The diagnostic output produced by the vision agent is therefore treated as a supporting finding rather than a standalone diagnosis, so that it can be reviewed alongside symptom, laboratory, and retrieved evidence before any result is presented to the user.

4.5 Diagnosis agent

This module works on the patient symptoms, laboratory reports and related medical findings. The diagnosis agent is intended to organize patient provided medical information and different relevant output from other agents. It combines symptoms, history, labs, and vision findings into ranked clinical hypotheses. Its evaluation should consider accuracy, completeness and consistency of its outputs using the predefined test set.

4.6 Risk assessment

This module is intended to examine the related risks using the information provided to it. This module gives the risk levels as low, moderate and high. According to the risk level the module suggests the medication to the related patient. High-emergency patients should be treated first, and healthcare specialists should know the patient's condition; this is why the module was built.

Potential areas include medication warnings, known drug interactions, contraindication-related information, and other risks covered by the implemented data sources and rules. A risk result should be presented as decision support, not as a confirmed clinical conclusion.

4.7 Backend integration

The backend integration layer serves as the single point of coordination for all processing modules, receiving validated requests from the user interface, dispatching them to the Vision, Drug Information, RAG, Diagnosis, and Risk Assessment modules, and aggregating their outputs into a unified

response payload for the results interface. The /ask endpoint also provides an interface for healthcare-related queries.

Integration testing is required to verify that the system works correctly. It should validate that each module connects to the other modules in the system and handles missing data, valid requests, invalid inputs and sources failure efficiently.

5 Implementation

5.1 Development environment and technology stack

The current backend of the project is implemented using Python and FastAPI. OpenFDA is used by the drug module to retrieve drug-label information. The RAG workflow includes PubMed and WHO health documents. The technologies used in the project are listed in Table 2.

Table 2. Technologies used in the project.

Component	Technology or resource	Purpose
Backend API	Python, FastAPI	Request handling and module integration
Drug information	OpenFDA	Retrieving available drug-label information
Literature source	PubMed	Medical literature retrieval
Health document source	WHO	Health-related reference documents
RAG processing	Python document-processing and retrieval modules	Preparing and retrieving evidence
Medical image analysis	Actual implemented vision framework and model	X-ray processing
User interface	Actual frontend framework	User input and result presentation
Data storage	Actual implemented storage mechanism	Storing supported records and outputs

5.2 Backend API

The implemented API includes the routes listed in Table 3.

Table 3. Implemented API routes.

Endpoint	Purpose
GET /	Check the API's basic availability
POST /ask	Submit a healthcare-related query
POST /drug/analyze	Request drug information analysis
POST /drug/interaction	Submit a drug interaction query
POST /report/upload	Upload a supported report
GET /reports	Retrieve the available report list

5.3 Prototype

The prototype is intended to provide the interface for users for medical relevant information. After login to the system, the user is directed to the main dashboard, where they can submit symptoms, upload X-ray images or laboratory reports, and request drug information or interaction checks through the interface that communicates with the backend routes.

5.4 Privacy and safety considerations

Healthcare data may contain patients' private information. Therefore, the system should give access to the authorized users only. Access to stored records and model outputs should additionally be logged and auditable, so that every retrieval, upload, and drug query can be traced back to the user who initiated it.

Medical findings must be presented with appropriate limitations. The system should not independently prescribe medication, modify treatment, or claim to replace a qualified healthcare professional.

6 Current status and evaluation plan

6.1 Current status

The current status of project development is given in Table 4.

Table 4. Current status of project components.

Component	Known status	Evidence or next step
FastAPI backend	Implemented	Verify available endpoints and responses
Drug analysis	Implemented at API level	Test medicine queries and source failures
OpenFDA integration	Available	Check retrieved fields and source coverage
Medicine-name normalisation	Partially developed	Test generic names, aliases, salts, and brand names
PubMed document processing	Implemented in part	Verify final article and chunk counts
WHO document processing	Implemented in part	Test retrieval against known queries
Drug interaction and risk assessment	Requires end-to-end verification	Test against reliable reference information
X-ray vision module	Separate project component	Confirm model, integration, and measured performance
Symptoms and laboratory module	Requires implementation-status confirmation	Verify actual analysis logic
Frontend integration	Requires end-to-end verification	Test navigation, requests, and results
Integrated system evaluation	Pending confirmation	Run repeatable functional and performance tests

6.2 Evaluation dataset and test cases

Evaluation should be conducted separately for each module before the complete system is assessed.

For the drug module, test cases should include recognised generic names, known aliases, unsupported names, and medicines for which the data source returns incomplete information.

For the RAG module, prepare queries with known relevant documents and assess whether the correct passages are retrieved. The final evaluation should record the number of test queries, retrieval settings, and criteria used to judge relevance.

For the vision module, use the actual labelled dataset and a held-out test set. Report the class distribution, model configuration, and metrics calculated from the test predictions.

For backend integration, test successful requests, invalid inputs, missing fields, unsupported file formats, and unavailable data sources.

6.3 Evaluation metrics

For different modules or agents different evaluation metrics are used.

- For RAG, evaluation may include retrieval precision, recall at a selected cutoff.
- For Drug module, evaluation may include accuracy for medicine name matching.
- For the complete application, evaluation may include API response time, successful request rate, error handling, and successful integration of the relevant modules.

No numerical values should be reported until the corresponding tests have been run.

7 Discussion

7.1 Integrated healthcare information

The system's modular design allows all the agents to work together and access through a common application. The practical value of this system depends on the reliability of each module and the correctness of the information exchanged between them.

7.2 Drug name and source coverage

The system is being developed for global use. Across the world, each medicine has different names, brand names and availability according to nation specifications. A global medicine-identification system therefore requires a carefully maintained mapping between recognized names and the corresponding active ingredients. OpenFDA is useful for this purpose, but it cannot be relied on as a worldwide medicine registry.

7.3 Human oversight

The system is developed to support healthcare specialists for decision making. This system is intended to augment not to replace the actual clinicians. After all, the final clinical decision remains with the responsible physician.

8 Limitations

8.1 Incomplete integration

The project is currently in the development phase, so individual modules are at different stages of development. Until the complete application is fully developed, integrated and tested, the system cannot be described as a fully validated end-to-end healthcare decision-support solution.

8.2 Data coverage

The drug-interaction and the RAG depend on the accuracy, coverage and the currency of sources collected.

8.3 Model validation

The performance of the medical image and other predictive modules must be established using suitable held-out data. Results from a small dataset or a development demonstration cannot establish clinical reliability.

8.4 Clinical use

The system is a project prototype and is not a substitute for professional medical advice, diagnosis, or treatment. Clinical deployment would require additional validation, privacy safeguards, and appropriate regulatory and ethical review.

9 Conclusion and future work

9.1 Conclusion

This paper describes the design and development of a Multi-Agent AI-Based Healthcare Decision Support System. The system is being built to bring together different healthcare functions in one platform, including drug information search, medical research retrieval, medical image analysis, and health risk assessment.

At present, the project includes a FastAPI-based backend, a drug analysis module connected to OpenFDA data, and document processing features that use information from PubMed and WHO resources. Additional modules are planned to support medical image analysis and other healthcare assessment functions in the future.

The main goal of this framework is to organize healthcare information in a structured and user-friendly workflow while clearly separating AI-generated insights from professional medical advice. The system's performance and reliability will be validated through both individual module testing and complete end-to-end testing. Any claims related to accuracy, effectiveness, or clinical benefits will only be made after proper evaluation and testing have been conducted.

9.2 Future work

Stage II will focus on integrating the Drug Information Agent, RAG module, Risk Assessment Module, Vision Agent, and Symptoms and Laboratory Module into a unified healthcare decision-support system. The drug module will be enhanced to recognise generic names, brand names, salt forms, and regional medicine names using verified medical databases and regulatory sources.

The RAG module will be improved by expanding PubMed and WHO document collections and providing relevant medical references with generated responses. The Vision Agent will be evaluated using fracture and chest X-ray datasets, with performance measured through accuracy, precision, recall, and F1-score.

The Risk Assessment Module will be enhanced to identify potential drug interactions, contraindications, and medication-related warnings. Further development will include integrating all modules through the FastAPI backend, improving the dashboard, and generating consolidated healthcare assessment reports.

Additional work will focus on expanding international medicine coverage, strengthening patient-data privacy, and evaluating the system's accuracy, response time, and reliability. The system will remain a decision-support tool, with final medical decisions made by qualified healthcare professionals.

References

- [1] Ouanes K, Farhah N. Effectiveness of artificial intelligence (AI) in clinical decision support systems and care delivery. *Journal of Medical Systems*. 2024;48(1):74. <https://doi.org/10.1007/s10916-024-02098-4>
- [2] Susanto AP, Lyell D, Widyantoro B, Berkovsky S, Magrabi F. Effects of machine learning-based clinical decision support systems on decision-making, care delivery, and patient outcomes: a scoping review. *Journal of the American Medical Informatics Association*. 2023;30(12):2050–63. <https://doi.org/10.1093/jamia/ocad180>
- [3] Kim Y, et al. MDAgents: an adaptive collaboration of LLMs for medical decision-making. In: *Advances in Neural Information Processing Systems*. Vol. 37. 2024. p. 79410–52.
- [4] Chen J, Lin H, Han X, Sun L. Benchmarking large language models in retrieval-augmented generation. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. 2024;38(16):17754–62. <https://doi.org/10.1609/aaai.v38i16.29728>
- [5] Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*. 2020;36(4):1234–40. <https://doi.org/10.1093/bioinformatics/btz682>
- [6] World Health Organization. *Medication without harm: WHO Global Patient Safety Challenge* [Internet]. Available from: <https://www.who.int/initiatives/medication-without-harm>
- [7] U.S. Food and Drug Administration. Drug product labeling API. openFDA [Internet]. Available from: <https://open.fda.gov/apis/drug/label/>
- [8] Çallı E, Sogancioglu E, van Leeuwen KG, Murphy K, van Ginneken B. Deep learning for chest X-ray analysis: a survey. *Medical Image Analysis*. 2021;72:102125. <https://doi.org/10.1016/j.media.2021.102125>
- [9] Öğdücü CU, Gürbüz S, Karaköse M, Hanoğlu E. Medical implications of LLM based clinical decision support systems in healthcare. In: *Proceedings of the 29th International Conference on Information Technology (IT)*. 2025. p. 1–4.
- [10] Ayers JW, et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. *JAMA Internal Medicine*. 2023;183(6):589–96.
- [11] Peek N, Capurro D, Rozova V, van der Veer SN. Bridging the gap: challenges and strategies for the implementation of artificial intelligence-based clinical decision support systems in clinical practice. *Yearbook of Medical Informatics*. 2024;33(1):103–14. <https://doi.org/10.1055/s-0044-1800729>
- [12] Maity S, Saikia MJ. Large language models in healthcare and medical applications: a review. *Bioengineering*. 2025;12(6):631. <https://doi.org/10.3390/bioengineering12060631>
- [13] Meng X, Yan X, Zhang K, Liu D, Cui X, Yang Y, et al. The application of large language models in medicine: a scoping review. *iScience*. 2024;27(5):109713. <https://doi.org/10.1016/j.isci.2024.109713>
- [14] Ong JCL, et al. Development and testing of a novel large language model-based clinical decision support systems for medication safety in 12 clinical specialties. *arXiv:2402.01741* [Preprint]. 2024. Available from: <https://arxiv.org/abs/2402.01741>

-
- [15] Achiam J, et al. GPT-4 technical report. *arXiv:2303.08774* [Preprint]. 2023. Available from: <https://arxiv.org/abs/2303.08774>
 - [16] Ji Z, et al. LLM internal states reveal hallucination risk faced with a query. *arXiv:2407.03282* [Preprint]. 2024. Available from: <https://arxiv.org/abs/2407.03282>
 - [17] Levra AG, et al. A large language model-based clinical decision support system for syncope recognition in the emergency department: a framework for clinical workflow integration. *European Journal of Internal Medicine*. 2025;131:113–20.
 - [18] Siddiky MNA, et al. *Optimizing AI language models: a study of ChatGPT-4 vs. ChatGPT-4o*. Preprints. 2025. <https://doi.org/10.20944/preprints202502.0066.v1>
 - [19] Sheikh AA, Chong EKP. Artificial intelligence-driven optimization for 3-D integrated circuit manufacturing: a system of systems framework. *IEEE Transactions on Components, Packaging and Manufacturing Technology*. 2025;15(7):1538–52.
 - [20] Nair M, Svedberg P, Larsson I, Nygren JM. A comprehensive overview of barriers and strategies for AI implementation in healthcare: mixed-method design. *PLOS ONE*. 2024;19(8):e0305949.