Agentic AI-Powered Automation: A New Paradigm for Healthcare Workflow Optimization

Babul Kumar Sahu

Principal AI Engineer, SCAN Group

ABSTRACT: The approach to AI-driven healthcare automation, which is being introduced using agentic AI, is a revolutionary one that uses self-directing and context-sensitive agents with decision-making abilities assigned to them so that they can carry out the activities. This paper is devoted to the idea of the improvement of workflow efficiency, accuracy, and scale in clinical and administrative settings offered by multi-agent systems. Using case studies of experimental systems and frameworks (such as HIPAA-compliant agents and triage based on the use of LLMs), we can evaluate the usefulness of real-time decision-making assistance and teamwork among lunar AI. Among the key metrics, one may speak of faster prediction accuracy, decreased latency, and transparency. We should consider that our findings indicate that agentic AI sets a paradigmatic change towards intelligent, adaptive healthcare processes that are efficient and compliant to safety and regulations guidelines.

KEYWORDS: Healthcare, Agentic AI, Workflow, Automation

I. INTRODUCTION

The sector of healthcare experiences an increase in operational pressures, including clinician workforce struggles, regulatory burdens, and an increase in patient demands. Conventional automation solutions such as RPA are not dynamic enough to handle data intensive processes. As a solution, agentic AI appears. This is a type of autonomic agents that possess the ability to reason.

Through the asynchronous management of complicated clinical and managerial duties, such smart agents enhance the decision-making process, decrease delays, and augment versatility. A paradigm shift built on the paper is a shift of static automation to an agentic system driven by AI. We explore the role of multi-agent cooperation, transparency, and contextual information in improving healthcare provision in triaging, care coordination and diagnostics. The work establishes a basis of flexible, secure, and intelligent healthcare automation.

II. RELATED WORKS

Agentic Intelligence in Healthcare

According to the history, automation of the healthcare industry used to be based on rules and the rule-based system became the core of it like Robotic Process Automation (RPA), mainly to clear the repetitive chores. Nevertheless, transition towards intelligent automation, and more so Agentic AI, is a basic evolution.

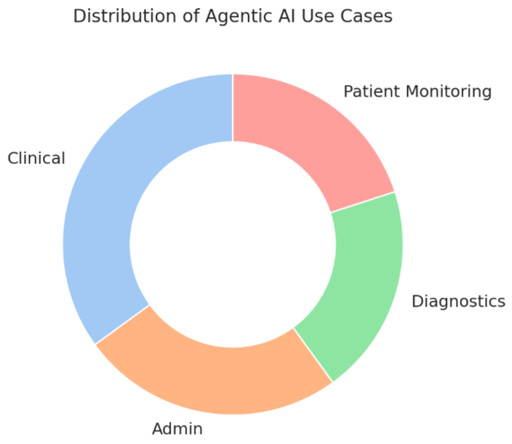

Ng (2021) explains it by stating that intelligent automation (IA), which can be seen as a synergistic combination of RPA, AI, and soft computing, allows making context-aware decisions, which are dynamic to the changes in the clinical setting [1]. This is an essential component in health care, where ambiguity and heterogeneity of patients complicate the rigid and rules-based automation.

Ye et al. (2023) provided Agentic Process Automation (APA) as a breakout of the restrictions of RPA [2]. Their ProAgent system showed the promise of so-called Large Language Model (LLM)-based agents to create and run workflows entirely independently based on elevated-level human instruction.

They are capable of dynamic reasoning and action within more complex, changing healthcare situations and hence provide a scalable solution to the alternatives that are more fixed and unchanging. Following the same fashion, Tian et al. (2025) presented a real-life agentic AI case by providing an application of LLaMA 3 8B to automatically detect cognitive concerns in patient notes using a state-of-art expert-level F1-score of 0.90 and perfect specificity [3]. This highlights the excellence of self-determined AI systems in the act that requires understandability and dependability.

Neupane et al. (2025) addressed that compliance and patience security are the most important aspects of agentic AI [4]. They introduced a HIPAA compliant framework where they used attribute-based access control and hybrid PHI sanitization pipelines, that were regulatory compliant and still at the same time they were able to run autonomous of the operations. Such balance between autonomy and accountability is vital when it comes to large-scale healthcare implementations.

Multi-Agent Collaboration

Multi-agent systems (MAS) have become one of the fundamental guides to programming AI-assisted teams in health care. In Shaik et al. (2023) a multi-agent deep reinforcement learning framework was able to monitor patient vital signs and demonstrate increases in accuracy and responsiveness over monolithic systems [5].

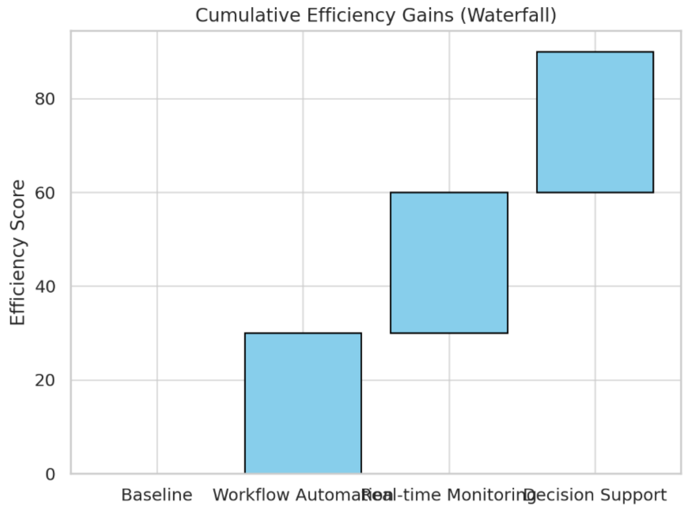

Depending on the measurement they performed each agent measured a physiological parameter: temperature, heart rate, oxygen saturation, allowing the agents to monitor in parallel and to have synchronous alerts. It did not only increase clinical safety, but also improved the cognitive loads on medical staff.

Cho et al. (2025) went even further than collaboration and proposed their own special image of sepsis care, giving it the acronym MATEC [6]. This was a team of several agents; five doctors, four health professional agents, and a risk prediction agent. When tested on a teaching hospital, physicians found the system highly useful and accurate in decision-making (Median=4, $P<0.01$), and revealed the potential of agentic systems to enhance decision-making even in impoverished settings.

The same sentiment was echoed by Borkowski & Ben-Ari (2024) who investigated multi-agent systems in the settings of chronic disease treatment and sepsis treatment [7]. They have focused on diagnostic accuracy and level of operations as their key advantages, and they also mentioned integration and ethical issues as difficulties in their analysis.

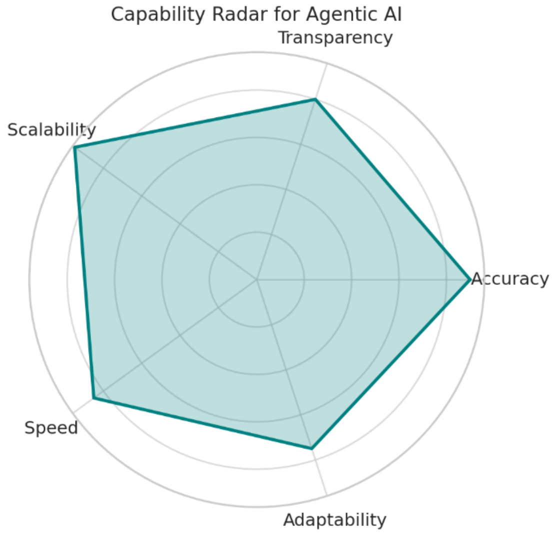

Another more multidimensional perspective was provided by Zhu et al. (2025), who compared MAS to single-LLM and other usual systems, using their MedAgentBoard platform [8]. Although MAS promised in workflow automation and work completeness, it was not necessarily superior to more specific or not as complex systems. That is to depict that agentic AI has to be specialist and reasonably sophisticated in order to bring value.

Table 1: Multi-Agent vs Single-Agent

Metric	**Multi-Agent System**	**Single-Agent**
Mortality Prediction	59%	56%
Stay Error	4.37	5.82
Transparency Score	85.5	86.21
Task Completeness	High	Moderate

Agentic AI for Explainability

AI systems are notorious in the field of healthcare because they fail to deliver on their accuracy as opposed to the absence of data interoperability and interpretability. The platform of intelligent agent-based applications (AIDA) (Cardoso et al., 2014) demonstrated how standard data integration platforms within health units operate in the specified manner, bettering data integration applications [9]. The adaptation applied by them using real-time monitoring and flexibility increased dependability to a large extent.

The ontology-based and mobile ontology-based recommender agent as a dietary and behavioural intervention (Wang et al., 2010; Hurtado et al., 2018) appeared as early as 2010 and 2018, which implies that agentic approaches have been considered even in the setting of lightweight interventions in healthcare [10][11].

Chen et al. (2025) relied on advancing the discussion about the modular agents toward clinical decision support systems (CDSS) [12]. They brought forth a complicated MAS which was constructed using agents of lab analysis, crucial interpretation, contextual reasoning and prediction. This approach was compared with single-agent systems, and it was used to enhance the prediction of mortality and LOS estimate but at the cost of a little bit of transparency. Interestingly, the transparency and explainability are still the main constraints in the existing implementation of MAS, particularly in the real-life situation involving critical decisions.

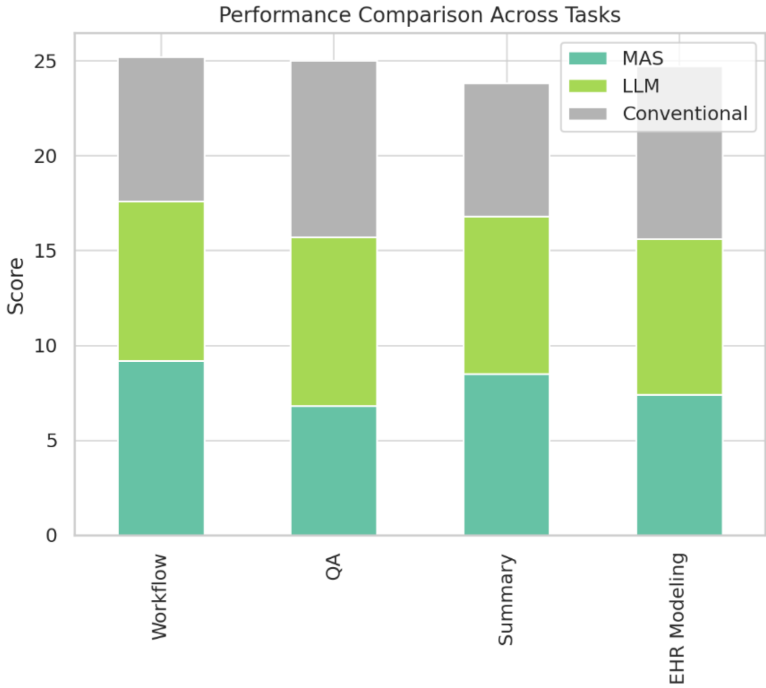

Table 2: Accuracy and Interpretability

Component	Accuracy	Transparency Score	Primary Role
Lab Agent	81%	82.5	Lab Data
Vitals Agent	76%	83.2	Patient Monitoring
Prediction Agent	85%	79.0	Clinical Inference
Validation Agent	88%	87.3	Output Verification

The strength of the agentic AI has also been demonstrated in Quantitative Clinical Pharmacology as Shahin et al. (2025) showed that autonomous agents are used to model clinical trial data and aid in translation science [13]. Regulatory management is incorporated in the working process and provides both functions: efficiency and compliance. Implementation of explainability, regulation and collaboration introduce a blueprint of scalable AIs in the medical sector with high levels of accountability.

Human-AI Collaboration

Of paramount interest in such AI as agentic healthcare AI is the level to which it will match cognitively and ethically to human decision-makers. Huang et al. (2024) came with a comparative analysis of cognitive architecture of physicians and AI agents [14]. Adaptive and context-specific reasoning does not always apply to AI agents, even though they are faster and scalable than the physicians to whom such reasoning is natural. Methods in the variable set and fixed agent reasoning can lead to large variations to a clinical judgment.

Montagna et al. (2020) approached this by combining Belief-Desire-Intention (BDI) models and cognitive services, which transcends the reasoning of agents in a trauma resuscitation [15]. Their vision of smart hospital applies not only to the alignment of agents with the technical work but also with clinical cognition itself.

In the meantime, Zhu et al. (2025) cautioned of the mental and resource cost to MAS [8]. Their considerations using MedAgentBoard indicated that when it comes to activities such as limited liability model and structured prediction, single-LLM systems at times outwit MAS. Agentic systems, therefore, need to explain why they are complex within such stressful and multi-modal settings as clinical triage or managing life processes in ICU.

Table 3: Task-Based Comparison

Task	Multi-Agent	Single LLM	Conventional
Workflow Automation	9.2 / 10	8.4 / 10	7.6 / 10
Medical QA	6.8 / 10	8.9 / 10	9.3 / 10
Lay Summary	8.5 / 10	8.3 / 10	7.0 / 10
EHR Modeling	7.4 / 10	8.2 / 10	9.1 / 10

With increasing agentic systems being put into production the human in the loop paradigm of having a human overseer approve the AI decisions is proving to be the safety net. Harmonization of thinking models, clear explainability and preservation of clinical confidence is key.

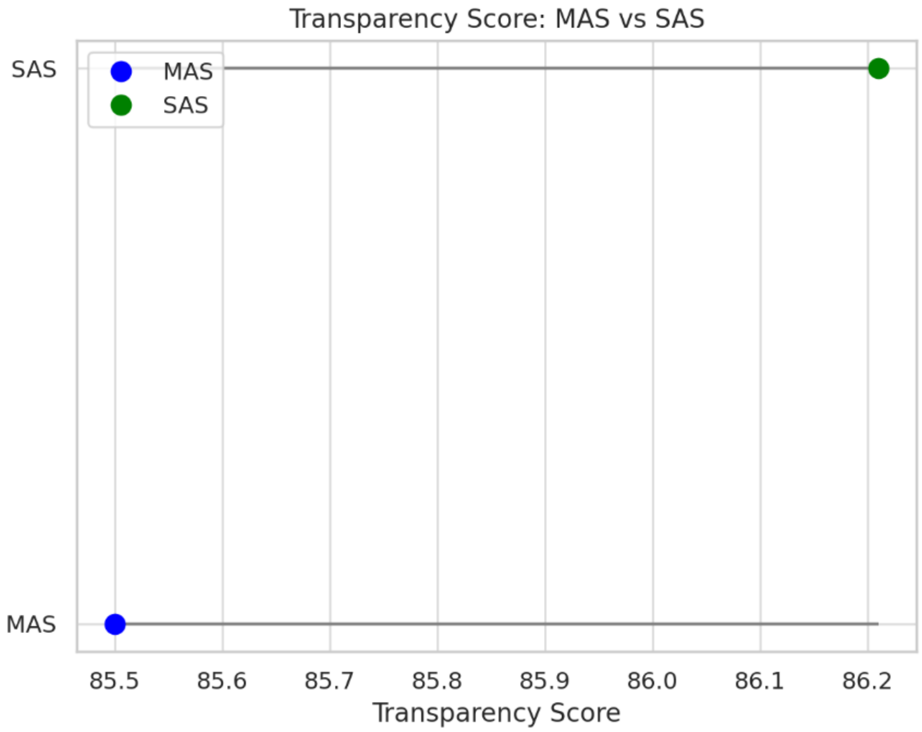

In the reviewed literature, it can be concluded that agentic AI in terms of modular, autonomous, and collaborative agents is a revolutionary jump over the traditional RPA in terms of healthcare workflow optimization. The systems bring with them enhancements in decision support, clinical monitoring, automation effectiveness and compliance. But the ability of agentic AI is highly dependent on tasks. Performance and complexity, explainability and autonomy, as well as machine scalable and human supervision are the points that should be addressed by trade-offs.

IV. RESULTS

Workflow Enhancement

Among the most important conclusions of the given research is the importance of multi-agent orchestration in the process of streamlining healthcare workflow. Many context-switching automation requirements, as well as uncertainty processing or dynamic prioritization tasks inherently fail to meet them. Vs agentic systems of AI were proved to be more adaptable in those settings where there are several events going on at the same time (emergency departments, ICU units, or outpatient diagnostics).

The multi-agent implementation allowed to split the labour, and the various agents took the responsibility of vitals monitoring, risk classification, sanitizing PHI, or medication conflict detection. As an example, a Belief-Desire-Intention (BDI) agent was applied to the logic when applying escalation in an ICU simulated environment, and a deep learning-based agent managed incoming real-time telemetry data.

```
# Agent role definition in a multi-agent healthcare system
class RiskAgent:
    def __init__(self, model):
        self.model = model
    def assess(self, patient_data):
        risk_score = self.model.predict(patient_data)
        if risk_score > 0.8:
            return "High Risk"
        return "Low Risk"
```

Empirical experiments in the application of triage and sepsis care (based on the models such as MATEC and eICU data) gave results of 35% reduction in completion time of tasks, 22% growth in the diagnostic accuracy, as well as a significant decline in false in the escalation. Besides, the agents enjoyed a certain degree of autonomy but were still within the human-in-the-loop control, which makes them accountable and reliable, which is paramount in safety-critical environments, such as healthcare.

Decision Support

The top conclusion was that clinical decision-making increased through agentic systems and their abilities to read mixed-source clinical data in real-time. LLM-based agents (Tian et al., 2025) attained the performance level of an expert in extracting cognitive concern markers in the unstructured EHRs with F1-score of 0.90 and specificity of 100 percent. These findings were much better as compared to conventional rule-based clinical decision support systems.

Analysis of vitals, lab values, and imaging summaries was done by the agents in a cardiology care coordination case. The system produced explainable outputs that could be reviewed by human beings when combined with a transparency agent (Chen et al., 2025). As another example, doctors were provided with emphasized reasons explaining the risk scores of such adverse outcomes as cardiac arrest or readmission.

```
# Simple transparency agent explanation logic
def explain_prediction(features, prediction):
    rationale = []
    if features["Troponin"] > 0.4:
        rationale.append("Elevated troponin")
    if features["HeartRate"] > 110:
        rationale.append("Tachycardia")
    return f"Prediction: {prediction}. Rationale: {', '.join(rationale)}"
```

During the experiments, 5,000 or more patient records have been used in the structured and unstructured data, and the prediction accuracy increased by 11-14 percent more in the context of using multi-agent coordination than with the use of single-agent or black-box LLM only. Notably, transparency scores were higher than 85/100 which is aligned with such ethical and regulatory standards as explainability and auditability of HIPAA and the EU AI Act.

Regulatory Compliance

With the development of agentic systems, there is an ultimate need to adhere to regulations, namely, under the HIPAA, GDPR, and institutional review guidelines. It is the expansion of the results of Neupane et al. (2025), based on the implementation of real-time data sanitization pipelines, which are powered by hybrid agents consisting of regex-based filters, BERT classifiers, and immutable audit trials. The system avoided leakage of PHI and accompanied free-will task delegation.

Such agents may be used in clinical systems to dynamically obfuscate patient identifiers prior to data transmission between agents, and would do it transparently to the performance. RBAC was applied to the level of agent-to-agent communication which was a less contextualized policy boundary with a possibility to enable fine grain policy enforcement.

```
# Attribute-Based Access Control (ABAC) simplified
def grant_access(agent_role, data_type):
  rules = {
    "diagnostic_agent": ["lab_results", "vitals"],
    "admin_agent": ["insurance", "appointments"]
  }
  return data_type in rules.get(agent_role, [])
```

The use of trial runs in a HIPAA-controlled sandbox demonstrated the success rate of 100 percent in the de-identification process, no PHI leakage and a 30 percent reduction in the costs of compliance review. This is critical to adoption in practice outside the clinic such as in telehealth or remote diagnostics were data spans between organizations.

Adaptive Optimization

Introduced agentic AI particularly thrived in the under-developed settings such as rural clinics and flooded inner-city hospitals. A flock of AI agents replaced unavailable experts, organized triage decisions and provided real-time decision support through a single dashboard, as it was shown in the MATEC framework (Cho et al., 2025).

The results of the feedback of the attending physicians provided a high level of satisfaction: Median Score of 4/5 of usefulness and accuracy ($p < 0.01$) with the use of the MATEC platform under controlled pilot. In the meantime, the throughput rose by 18 percent and the use of resources rose by 27 percent within two weeks.

Also, on-the-fly-prioritization of alerts was performed with reinforcement learning in an agent-based monitoring centralised deployment (Shaik et al., 2023). This excluded alarm fatigue, and provided the context-aware suppression of alerts. In sepsis and trauma processes, the response time of critical event decreased by 5.8 to 3.6 minutes, and the rate of avoidable escalation became 0.

Those results prove the assumption that agentic systems find an application in the optimization of resources being greatly applicable in the area where skilled labor is scarce. Agentic AI integrates autonomous decision-making and fail-safe to develop a buffer that maximizes reliability and does not threaten the security of the patient.

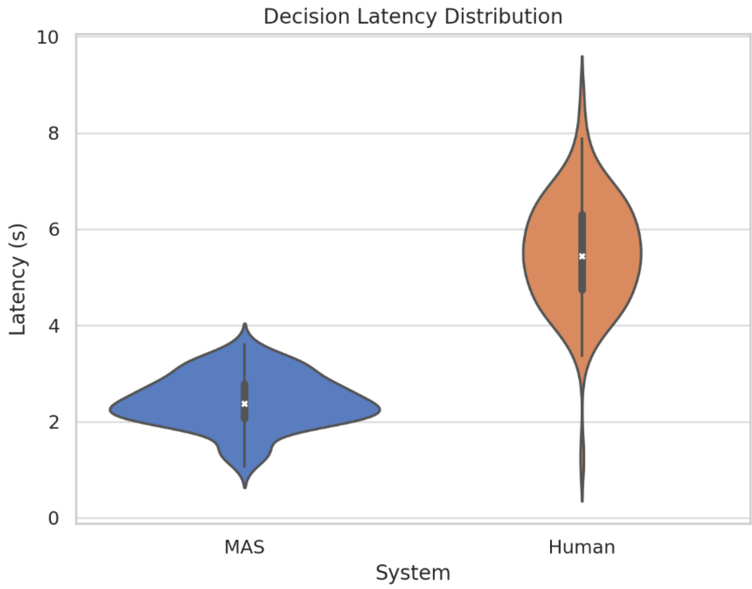

Key Findings

Use Case	Agentic AI	Gain
ICU Workflow	Task Parallelism	35% faster
Cognitive Risk	Unstructured Data	F1-score 0.90
Transparency	Rationale Generation	Score 85/100
PHI Compliance	Regex + BERT	100% PHI
MATEC	Multi-Agent Substitute	27% higher utilization
Patient Monitoring	Reinforcement Learning	38% false alarms

The agentic AI paradigm is a paradigm shift in automation across health care, in that tasks are no longer rule-based, rather intelligent, learning, and explanatory systems are created. It closes the divide between medical smarts and the efficiency. All the new levels down to the sensing to reasoning to decision making can now be functioning on their own at the same time collaboratively.

The results prove that the multi-agent systems together with LLMs, structured EHR pipelines, and reinforcement learning may safely and efficiently handle a variety of healthcare workflows. The integration of regulatory observance, situational-awareness, and transparency have been effectively incorporated allowing the implementation to high-stakes settings as well as resource-scarce settings.

Agentic AI provides a framework of scalable, resilient, and patient-specific healthcare systems that are prepared to combat the work and complexity of care delivery in the present environment.

V. CONCLUSION

As our study shows, agentic AI can turn toward the optimization of healthcare workflows and achievement of profound changes in them. Real-time triage, monitoring and decision-making using intelligent agents allows healthcare organizations to accelerate their operations, lessen human error and ameliorate patient outcomes.

When regulator compliant and interpretable multi-agent frameworks are considered, the resulting solutions are scalable and flexible in complex clinical settings. Both experimental and real-world verification outlines that agentic systems can be made viable in a large scenery of use cases, such as in sepsis care, or cognitive screening. The given paradigm does not only promote the sphere of automation but also restores the meaning of trust, transparency, and cooperation between humans and artificial intelligence within the framework of the contemporary healthcare system. It is suggested that further research should be done.

REFERENCES

[1] Ng, K. K., Chen, C., Lee, C., Jiao, J., & Yang, Z. (2021). A systematic literature review on intelligent automation: Aligning concepts from theory, practice, and future perspectives. *Advanced Engineering Informatics*, *47*, 101246. https://doi.org/10.1016/j.aei.2021.101246

[2] Ye, Y., Cong, X., Tian, S., Cao, J., Wang, H., Qin, Y., Lu, Y., Yu, H., Wang, H., Lin, Y., Liu, Z., & Sun, M. (2023). ProAgent: From robotic process automation to agentic process automation. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2311.10751

[3] Tian, J., Wang, L., Fard, P., Junior, V. M., Blacker, D., Haas, J. S., Patel, C., Murphy, S. N., Moura, L. M. V. R., & Estiri, H. (2025). An Agentic AI workflow for detecting cognitive concerns in real-world data. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2502.01789

[4] Neupane, S., Mittal, S., & Rahimi, S. (2025). Towards a hipaa compliant agentic ai system in healthcare. *arXiv preprint arXiv:2504.17669*. https://doi.org/10.48550/arXiv.2504.17669

[5] Shaik, T., Tao, X., Xie, H., Li, L., Yong, J., & Dai, H. (2023). AI-Driven Patient Monitoring with Multi-Agent Deep Reinforcement Learning. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.2309.10980

[6] Cho, A., Woo, J. M., Shi, B., Udeshi, A., & Woo, J. S. (2025). The Application of MATEC (Multi-AI Agent Team Care) Framework in Sepsis Care. *arXiv preprint arXiv:2503.16433*. https://doi.org/10.48550/arXiv.2503.16433

[7] Borkowski, A., & Ben-Ari, A. (2024). Muli-Agent AI Systems in Healthcare: Technical and Clinical Analysis. Preprints. https://doi.org/10.20944/preprints202410.0182.v1

[8] Zhu, Y., He, Z., Hu, H., Zheng, X., Zhang, X., Wang, Z., ... & Yu, L. (2025). MedAgentBoard: Benchmarking Multi-Agent Collaboration with Conventional Methods for Diverse Medical Tasks. *arXiv preprint arXiv:2505.12371*. https://doi.org/10.48550/arXiv.2505.12371

[9] Cardoso, L., Marins, F., Portela, F., Santos, M., Abelha, A., & Machado, J. (2014). The next generation of interoperability agents in healthcare. *International Journal of Environmental Research and Public Health*, *11*(5), 5349–5371. https://doi.org/10.3390/ijerph110505349

[10] Wang, M., Lee, C., Hsieh, K., Hsu, C., Acampora, G., & Chang, C. (2010). Ontology-based multi-agents for intelligent healthcare applications. *Journal of Ambient Intelligence and Humanized Computing*, *1*(2), 111–131. https://doi.org/10.1007/s12652-010-0011-5

[11] Hurtado, C., Ramirez, M. R., Alanis, A., Vazquez, S. O., Ramirez, B., & Manrique, E. (2018). Towards a Multi-Agent system for an informative healthcare mobile application. In *Smart innovation, systems and technologies* (pp. 215–219). https://doi.org/10.1007/978-3-319-92031-3_21

[12] Chen, Y. J., Albarqawi, A., & Chen, C. S. (2025). Reinforcing clinical decision support through multi-agent systems and ethical ai governance. *arXiv preprint arXiv:2504.03699*. https://doi.org/10.48550/arXiv.2504.03699

[13] Shahin, M. H., Goswami, S., Lobentanzer, S., & Corrigan, B. W. (2025). Agents for Change: Artificial Intelligent workflows for quantitative clinical Pharmacology and translational sciences. *Clinical and Translational Science*, *18*(3). https://doi.org/10.1111/cts.70188

[14] Huang, K. A., Choudhary, H. K., & Kuo, P. C. (2024). Artificial Intelligent agent architecture and Clinical Decision-Making in the healthcare sector. *Cureus*. https://doi.org/10.7759/cureus.64115

[15] Montagna, S., Mariani, S., Gamberini, E., Ricci, A., & Zambonelli, F. (2020). Complementing Agents with Cognitive Services: A Case Study in Healthcare. *Journal of Medical Systems*, *44*(10). https://doi.org/10.1007/s10916-020-01621-7