

Article

Not peer-reviewed version

---

# A Survey of LLM-based Multi-agent Systems in Medicine

---

[Yanna Lin](#), Shaojie Xu, Wenshuo Zhang, Yushi Sun, Zixin Chen, Yanjie Zhang, Rui Sheng\*

Posted Date: 20 October 2025

doi: 10.20944/preprints202510.1448.v1

Keywords: multi-agent system; survey; medical domain



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# A Survey of LLM-Based Multi-Agent Systems in Medicin

Yanna Lin<sup>1</sup>, Shaojie Xu<sup>2</sup>, Wenshuo Zhang<sup>1</sup>, Yushi Sun<sup>1</sup>, Zixin Chen<sup>1</sup>, Yanjie Zhang<sup>1</sup>, Rui Sheng<sup>1,\*</sup>

<sup>1</sup> Hong Kong University of Science and Technology

<sup>2</sup> Huazhong University of Science and Technology

\* Correspondence: rshengac@connect.ust.hk

## Abstract

Large Language Model (LLM)-based multi-agent systems have shown great potential in supporting complex tasks in the medical domain, such as improving diagnostic accuracy and facilitating multidisciplinary collaboration. However, despite the advancement, there is a lack of structured frameworks to guide the design of these systems in medical problem-solving. In this paper, we conduct a comprehensive survey of existing medical multi-agent systems, and propose a medical-specific taxonomy along three key dimensions: team composition, medical knowledge augmentation, and agent interaction. We further outline several future research directions, such as incorporating human-AI collaboration to ensure that human expertise and multi-agent reasoning jointly address complex clinical tasks, designing and evaluating agent profiles, and developing self-evolving systems that adapt to evolving medical knowledge and rapidly changing clinical environments. In summary, our work provides a structured overview of medical multi-agent systems and highlights key opportunities to advance their research and practical deployment.

**Keywords:** multi-agent system; survey; medical domain

## 1. Introduction

Large Language Model (LLM)-based multi-agent systems have recently gained significant attention for their potential to support complex decision-making processes in the medical domain. By leveraging the complementary reasoning and collaboration of multiple agents, such systems aim to address the limitations of single-agent approaches, including hallucination, lack of domain specialization, and difficulties in handling multi-step reasoning. Recent studies have explored their use in diverse medical tasks, such as clinical diagnosis [1,2], clinical triage [3], and clinical trial design and optimization [4]. These applications highlight the promise of multi-agent systems in improving reliability, interpretability, and scalability in healthcare-related AI solutions.

Despite these advances, designing effective multi-agent systems for medicine remains highly challenging. To address this, an increasing number of studies have proposed various approaches, such as optimizing agent role allocation and collaboration strategies [5,6], incorporating domain knowledge into agent communication [2], and designing mechanisms to organize agents' discussion processes [7]. For example, MDAGent introduces a framework that dynamically determines whether LLM-based agents should work individually or collaboratively according to the complexity of medical tasks, mirroring real-world medical decision-making [5]. As the number of these efforts continues to grow, they remain fragmented, underscoring the need for a systematic understanding to help researchers clearly grasp the current landscape of these approaches and highlight future opportunities for designing multi-agent systems in medicine.

Several survey papers have begun to examine the development of multi-agent systems. However, many are domain-agnostic and therefore overlook medical-specific characteristics, such as designing agents to reflect the clinical workflows and physician specializations [8,9]. Some only focus on investigating the use of multi-agent systems for hospital simulation like hospital operations [9,10]. A few

surveys do discuss medical applications, but they either emphasize usage scenarios [11] or do not highlight the unique aspects brought by multiple agents, such as their interactions [12]. Moreover, existing general taxonomies—for example, interaction modes like centralized, decentralized, hierarchical, and shared message pool [8]—are also insufficient for medical contexts, as they overlook domain-specific mechanisms in clinical tasks, such as multidisciplinary team (MDT)-style collaboration or stage-wise clinical task allocation. To our knowledge, no prior survey has systematically structured existing work from this design-oriented perspective, nor provided a medical-specific coding framework that better reflects how multi-agent systems can be designed and deployed in real-world medical problem-solving tasks.

To fill this gap, our survey takes a design-oriented perspective on medical multi-agent systems. Specifically, we target to answer three questions: (1) *how to compose a multi-agent team to address practical problems*, (2) *how to empower agents with medical knowledge*, and (3) *how agents interact with each other*. To address these questions, we collected 50 papers that specifically focus on using LLM-based multi-agent systems for real-world medical problem-solving. We first analyzed these works and derived a medical-specific taxonomy along three dimensions: team composition, knowledge enhancement, and interactions, as shown in the Appendix A (Table 1). We further discuss broader challenges in current approaches and outline promising directions for the future development of medical multi-agent systems. These include incorporating human–AI collaboration to ensure that human expertise and multi-agent reasoning jointly address complex clinical tasks; enabling safe and transparent self-evolution in response to new medical knowledge and rapidly changing clinical environments; achieving deeper multimodal integration for richer cross-modal reasoning; designing and evaluating agent profiles that balance realism with adaptability; and expanding to more diverse clinical and public health scenarios. We hope that our work serves as a starting point for researchers and practitioners to better understand and design next-generation medical multi-agent systems.

## 2. Taxonomy of Medical Multi-Agent Systems

To build a comprehensive paper corpus, we first conducted a keyword search in Google Scholar using the query: *(LLMs OR large language models) AND (multi-agent OR multiple agents) AND (medicine OR medical OR clinical OR diagnosis)* and collected 15 seed papers. Starting with these seed papers, we iteratively expanded the corpus by tracing references. This process continued until saturation, when no new relevant papers were found, resulting in a total of 64 papers. Next, two authors independently reviewed the abstracts and introductions to exclude works unrelated to medical problem-solving (e.g., general-purpose systems [13], hospital simulations [14,15], or purely biological or genomic research [16]). This filtering yielded 50 papers for the subsequent analysis. We then analyzed these papers along three design dimensions of multi-agent systems: team composition, medical knowledge enhancement, and agent interaction. Each paper was independently coded by at least two authors. Note that, at this stage, we did not have unified the terms for all the coded perspectives. Subsequently, all authors convened for a series of weekly meetings dedicated to deliberating and refining the coding outcomes until a consensus was established. This process resulted in a final taxonomy that captures the key design patterns of medical multi-agent systems, as shown in Table 1. In the following, we will present this taxonomy across three aspects one by one: team composition, medical knowledge augmentation, and agent interaction.

**Table 1.** Taxonomy of LLM-based Multi-agent Systems in Medicine

Category	Subcategory	Sub-subcategory	Related Work
Team Composition		Clinical task allocation	[4,6,17–28]
		Specialization-oriented assignment	[5,6,19,20,22,23,29–41]
		Process-oriented allocation	[7,28,32,33,38,40,42–59]
		Expertise-level assignment	[2,35,47,56,60]
		Automatic assignment	[7,23,24,31–36,38,48,56,59]
Medical Knowledge	Agent-intrinsic	Role-play prompting	[3,5–7,23,24,29–34,37,39,45,56,60]
		Pre-trained model utilization	[17,20,22,28,38,44,57]
		Model fine-tuning	[2,6,18,20–22,28–30,50,58]
	Externally-assisted	Traditional medicine tool utilization	[34,37,42,43,52]
		Domain-specific model calling	[2,4,19,22,34,38,40,50,52,57]
		Medical knowledge-based RAG	[2–4,17–19,21,25,26,33,38–41,44,48,50–55,57–59]
Agent Interaction	Hierarchical Coordination	Agent recruitment	[6,19,27,31,32,34,41,47,48,56,58]
		Task delegation	[4,19,22,48,56]
		Joint discussion	[2,3,22,32,45,56,59]
		Information integration	[2–6,19,22,28–32,34–38,40,41,43,44,48,54–60]
	Peer Collaboration	Collaborative Discussion	[2,5,7,23–25,28,29,31,33–35,38,39,42,45,56]
		Clinical task-driven flow	[2,6,17,18,21–24,26–28,32,34,48,60]
		Problem-solving cycles	[3,19,20,22,25,30,32,38,40,41,46–59]

### 2.1. Team Composition

Team composition determines how agent roles are configured to address medical tasks. We identify five approaches to configuring the roles of agents in existing medical multi-agent systems.

**Clinical task allocation:** Agent roles are defined according to specific clinical tasks (e.g., diagnosis, prognosis, treatment planning) based on task requirements. Clinical task allocation emphasizes the organization of agent roles around specific medical tasks. Clearly defined task boundaries provide a foundation for subsequent task-level performance monitoring and behavioral pattern analysis. A representative example is presented by [26], who developed a multi-agent system for managing inpatient pathways. In their framework, agents such as the Admission Agent, Diagnosis Agent, Treatment Agent, and Discharge Agent are each responsible for distinct clinical phases, collaboratively covering the full patient journey from admission to discharge. Clinical task allocation is also applicable to finer-grained task decomposition in system design. For instance, [17] designed a three-agent system where agents were responsible for the overall condition assessment, antibiotic recommendation, and compliance checking against clinical guidelines.

**Specialization-oriented assignment:** Agents roles are aligned with distinct medical specialties (e.g., radiology, pathology, pharmacology). This mirrors the role structures in real-world hospitals, enhancing diagnostic accuracy and ensuring strong system scalability. For instance, [29] introduce a system where a General Practitioner performs triage and refers patients to a team of expert agents, each handling domain-specific tasks such as imaging interpretation or information synthesis. A Director agent then coordinates the discussion and generates the final diagnostic report. Similarly, [6] use a Triage Doctor to route cases to specialists, whose opinions are integrated by an Attending Physician. This structure effectively simulates multidisciplinary teams (MDTs), where specialists collaborate to improve care for complex cases. [30] design an MDT-inspired agent team for Alzheimer’s diagnosis, where agents such as the Primary Care Physician, Neurologist, Psychiatrist, and Geriatrician each focus on complementary aspects of patient assessment. Their findings are synthesized by an AD Specialist agent to produce a final risk evaluation. This design enables agent-level simulation of MDT collaboration and enhances diagnostic performance in complex cognitive disorders.

**Process-oriented allocation:** Agent roles are defined based on stages of the decision-making or task completion workflow, such as planning, analysis, refine, and final decision-making. Compared to clinical task allocation and specialization-oriented assignment, which emphasize real-world clinical practices and role-specific responsibilities, process-oriented allocation assigns conceptual task flows to agents. By leveraging abstract cognitive structures to restructure problem-solving pathways, it transcends the constraints of conventional clinical thinking and enables more innovative and systematic forms of intelligent collaboration. For example, in the “Generation—Verification—Reasoning” task flow proposed by [43], the Generator agent generates preliminary diagnostic or treatment plans based on predefined argumentation templates; the Verifier agent challenges these plans by posing structured critical questions, prompting the system to generate rebuttals or alternative arguments; finally, the Reasoner agent synthesizes the discussion and arrives at an acceptable decision. Similarly,

[46] designed three corresponding agents based on the “Generation—Evaluation—Optimization” task flow, responsible for generating initial proposals, evaluating the proposals, and optimizing them. This approach can be used to ensure the reliability and safety of medical decisions incorporating roles for review, feedback, and refinement.

**Expertise-level assignment:** Agent roles reflect different expertise levels, such as junior vs. senior physician roles. For instance, [60] implemented a multi-tiered diagnostic review system in which Junior Resident I was responsible for the initial diagnosis, followed by Resident II who critically evaluated the diagnosis from a peer-review perspective, aiming to identify cognitive biases such as anchoring bias and confirmation bias. A Senior Physician played a supervisory role, identifying and correcting cognitive distortions, and offering guidance and decision-making support. [47] proposed a risk-aware routing mechanism for the delegation of surgical error detection tasks across three professional tiers: the Resident-level, Attending-level, and Expert-level. Agents at the resident level employed checklist-based conservative reasoning; attending-level agents integrated structured evaluations with contextual interpretations; and expert-level agents incorporated multi-scale temporal pattern recognition to provide high-level insights. Expertise-level assignment simulates multi-tiered collaboration among physicians of varying seniority, effectively enhancing the depth of decision review and the correction of cognitive biases.

**Automatic assignment:** Agent roles are automatically defined or selected by algorithms or optimization strategies that generate or select the most suitable agents for a given medical task. Dynamically selecting optimal agents through algorithmic strategies significantly enhances the system’s adaptability and generalization across diverse task scenarios. For instance, [29] constructed a domain-specific expertise table for various LLMs, systematically quantifying each model’s strengths across medical domains, which enables the system to recruit an optimal subset of agents with demonstrated proficiency in the relevant subject areas and query difficulties. [36] proposed the Rotation Agent Collaboration (RAC) mechanism, in which a leading agent is dynamically selected based on the inferred intent of the question. This agent gathers information through polling from other agents and, after fusing the responses, designates the most suitable agent to make the final decision. [48] adopted a strategy of dynamically selecting AssistAgents aligned with the medical domains relevant to each query, assigning each agent to retrieve and synthesize evidence within its area of expertise. For medical question answering, [7] generated domain-specific agents based on the domains associated with both the question and the answer options. To address challenges in rare disease diagnosis and treatment, [34] designed an Attending Physician Agent that selects the most relevant specialists from a predefined pool based on the patient’s clinical profile and forms a MDT to reach diagnostic consensus.

## 2.2. Medical Knowledge Augmentation

Equipping agents with medical knowledge is essential to ensure reliable reasoning and improve task accuracy in clinical contexts. Existing approaches can be broadly grouped into two categories: *agent-intrinsic* methods, which enhance knowledge within the agent, and *externally-assisted* methods, which integrate external knowledge sources to agents but not modify the agent models.

### 2.2.1. Agent-Intrinsic

Agent-intrinsic methods enhance the medical knowledge embedded within the agents themselves. These approaches represent a spectrum of increasing specialization, ranging from lightweight prompt engineering to more intensive modifications of the model’s underlying parameters. This progression allows for tailored enhancement of an agent’s expertise, moving from broad role simulation to deep, task-specific knowledge integration.

**Role-play prompting:** Agents are guided by assigning them specific medical roles (e.g., cardiologist, nurse, patient) through carefully crafted prompts. By framing the task within a professional persona, the LLM is encouraged to adopt a communication style, reasoning process, and knowledge domain relevant to that role. This is a zero-shot, computationally efficient method to improve the quality and relevance of the agent’s output without altering the base model. For instance, several

frameworks simulate multi-disciplinary team (MDT) consultations via role-play [37,45]. A typical example is MedAgents [23], which introduces a framework for addressing domain-specific terminology and expert reasoning in the medical field using large language models. By enabling the model to “role-play” different medical experts, MedAgents facilitates multi-turn collaborative discussions to analyze and solve medical problems, thereby improving reasoning accuracy and interpretability.

**Pre-trained model utilization:** Agents leverage LLMs pre-trained on large-scale and curated medical datasets, instead of general-purpose LLMs. This approach overcomes the inherent knowledge limitations of generalist models by providing medical knowledge, such as complex medical terminology, clinical concepts, and reasoning patterns. For example, WSI-Agents [38] leverages a “MLLM model library” comprising five pre-trained multimodal large language models (e.g., WSI-LLaVA [61] and Quilt-LLaVA [62]) specialized for Whole Slide Image (WSI) analysis. Other systems also rely on medically-informed base models, such as CardAIC-Agents [22] employs MedGemma [63] as the foundation for its multidisciplinary discussion tools.

**Model fine-tuning:** As the most intensive method, fine-tuning adapts models to highly specific medical tasks or datasets by further training them. This process adjusts the model’s weights, enabling it to master specialized knowledge, adhere to specific clinical guidelines, or adopt a particular reporting style. It overcomes the generic nature of pre-trained models by instilling deep, task-specific expertise. For example, the agents in MMedAgent-RL [6] such as the “triage doctor” and “attending physician” are fine-tuned by using GRPO [64] (a kind of RL method) to impart domain knowledge and optimize collaborative policies and decision-making strategies based on feedback. MRGAges [20] fine-tunes base BioMedGPT [65] models for agents on dedicated disease-specific subsets in IU X-ray [66] and MIMIC-CXR [67] to improve medical report generation.

### 2.2.2. Externally-Assisted

While agent-intrinsic methods enhance the inherent capabilities of LLMs, relying solely on the internal knowledge of these models presents significant challenges in the medical domain. LLMs are prone to factual hallucinations [68], lack the ability to process specialized data formats (e.g., genomic VCF files, ECG signals), and cannot execute deterministic computations required by many clinical protocols. To overcome these limitations, externally-assisted approaches equip agents with the ability to call upon external tools, models, and knowledge bases. This paradigm allows the LLM to function as a central orchestrator or “brain”, coordinating which external knowledge sources or tools to invoke.

This approach is exemplified by complex, hybrid systems that integrate multiple forms of external assistance. For instance, the DeepRare system [19] is designed for rare disease diagnosis by using an LLM as a central host that coordinates several specialized agents. These agents, in turn, invoke a variety of external medical tools and databases to perform evidence retrieval and diagnostic reasoning. Such systems illustrate a spectrum of external augmentation, which can be categorized by the increasing level of intelligence and complexity of the external resource, progressing from deterministic tools to specialized predictive models, and finally to dynamic knowledge retrieval systems.

**Traditional medicine tool utilization:** At the most fundamental level, agents employ established, often deterministic, medical tools as auxiliary supports. These tools include PubMed search engines, EHR retrieval systems, and clinical calculators. Their integration is critical for tasks requiring high fidelity, procedural accuracy, and the processing of structured or non-textual data. By offloading these functions, agents can ground their reasoning in reliable, standardized outputs, overcoming the non-deterministic and interpretive nature of LLMs. For example, a genotype analysis agent in DeepRare [19] uses Exomiser [69] (a specialized bioinformatics tool) to perform variant annotation and prioritization based on criteria like predicted pathogenicity, allele frequency, and genetic inheritance patterns. This process yields a precise list of candidate pathogenic genes, achieving a level of diagnostic accuracy in genomics that is unattainable for a generalist LLM. Similarly, the CardAIC-Agents framework [22] features a “CardiacExperts Agent” that utilizes NeuroKit2 (a Python toolbox) for processing raw ECG signals to obtain 12-leads ECG measurements.

**Domain-specific model calling:** A step beyond traditional tools, this approach involves agents integrating specialized AI models, such as radiology image classifiers or drug-disease interaction predictors. These models, often based on deep learning, provide sophisticated pattern recognition and predictive capabilities that complement the broad reasoning of an LLM. Calling these models allows the agent system to leverage deep, task-specific expertise learned from vast amounts of data, enabling more accurate analysis in domains like medical imaging or computational biology. For instance, DeepRare [19] calls upon PhenoBrain [70], a model that performs analysis on structured Human Phenotype Ontology (HPO) terms. Using classical machine learning and ontology matching, PhenoBrain rapidly generates an interpretable, probability-scored list of candidate diseases, enhancing the efficiency of the phenotype-driven diagnostic process by providing a quick and explainable differential diagnosis list for the LLM to reason upon. The CardiacExperts Agent in ardAIC-Agents [22] invokes multiple specialized models, including a fine-tuned multimodal cardiac diagnosis model considering lab data, ECG, and ultrasound, along with view classification and segmentation models for analyzing medical images.

**Medical knowledge-based RAG:** The most advanced form of external assistance involves Retrieval-Augmented Generation (RAG), where agents dynamically query knowledge bases and incorporate the retrieved information into their reasoning process at inference time. This method directly mitigates LLM hallucinations by grounding responses in factual, up-to-date information from structured (e.g., knowledge graphs) or unstructured (e.g., clinical guidelines, biomedical literature) sources. This ensures that the agent's outputs are not only contextually relevant but also transparent and traceable, fostering greater clinical trust. For example, DeepRare [19], employs a multi-faceted RAG strategy. Its knowledge retrieval agent implements an LLM-based RAG workflow to query medical knowledge bases like PubMed, Orphanet, and OMIM. The retrieved text is then summarized by a lightweight LLM to generate a transparent, traceable evidence chain that explicitly links diagnostic conclusions to source material. KERAP [53] uses a Retrieval Agent to extract and summarize entity relationships from a knowledge graph to inform diagnosis.

### 2.3. Agent Interaction

Interaction mechanisms define how agents coordinate with each other to accomplish medical tasks. Effective interaction is critical for balancing efficiency and reliability in clinical contexts. We categorize existing approaches into two broad types: *hierarchical coordination*, where agents are organized in a layered structure (such as task delegators or aggregators that coordinate and integrate the work of other agents), and *peer collaboration*, where agents interact on a more equal footing.

#### 2.3.1. Hierarchical Coordination

**Agent recruitment:** In hierarchical settings, higher-level agents are responsible for instantiating or selecting subordinate agents with specific expertise to address a clinical scenario. This mechanism allows the system to dynamically assemble a team with the necessary domain knowledge. For example, an upper-level controller may recruit genetic and cardiovascular agents when facing a case involving both hereditary and symptomatic factors [13]. Similarly, in the MMedAgent framework, a triage doctor first analyzes multimodal patient inputs to determine the appropriate specialty, and only the corresponding specialist agents (e.g., radiologists for X-ray images) are recruited, while others remain inactive [6]. This targeted activation ensures efficient use of resources while still providing specialized reasoning for the clinical case.

**Task delegation:** Higher-level agents distribute subtasks to subordinate agents, thereby structuring the workflow into well-defined stages. This delegation enforces a clear division of labor: the main agent generates an initial diagnosis or hypothesis, then assigns targeted subtasks (e.g., evidence retrieval, calibration, or domain-specific reasoning) to domain experts. Subordinate agents return their findings, which are then aggregated by the higher-level agent [48]. Such delegation improves efficiency by parallelizing subtasks while maintaining control at the supervisory level. For example, in a forensic investigation of a male fatality caused by aortic dissection after a car accident, the planner decomposes

the case into subtasks such as analyzing rupture characteristics, linking trauma with pre-existing conditions (e.g., hypertension, coronary heart disease), ruling out poisoning, and differentiating rib fracture causes. These subtasks are then distributed to specialized solvers (e.g., autopsy analyzer, medical history integrator, toxicology interpreter, trauma classifier), each returning results that the planner synthesizes into the final conclusion [71].

**Joint discussion:** In addition to strict delegation, hierarchical systems may incorporate guided group deliberation, where higher-level agents act as moderators or evaluators of subordinate discussions. Usually, rather than participating as equals, subordinate agents can contribute candidate solutions, while the higher-level agent comments, provides feedback, and ensures alignment with the overarching diagnostic goal. This structure balances open exchange with centralized oversight, ensuring that junior agents' opinions are considered without compromising clinical reliability. For example, in a clinical setting, junior residents propose and critique preliminary diagnoses, while a senior doctor moderates the exchange, identifies cognitive biases, and steers the group toward a refined conclusion, supported by a recorder who consolidates outcomes [45].

**Information integration:** Finally, hierarchical coordination culminates in the synthesis of subordinate outputs into a unified decision or recommendation. The higher-level agent integrates intermediate results, weighing the evidence and resolving conflicts across subordinate reports. This step is crucial in clinical contexts, as it ensures that diverse sources of reasoning—such as genetic evidence, clinical manifestations, and patient history—are combined into a coherent and trustworthy medical conclusion [58]. By maintaining a supervisory “review–integrate–decide” cycle, hierarchical systems promote both interpretability and accountability. For example, in the MAM framework, the diagnostic process is decomposed into multiple specialized roles—including general practitioners, specialist teams, radiologists, medical assistants, and a chief physician—each embodied by an LLM-based agent. The chief physician coordinates the discussion, synthesizes opinions and retrieved evidence into interim reports, and the specialist group votes on whether to endorse these reports. Once consensus is reached, the chief physician consolidates the results into the final diagnosis [29].

### 2.3.2. Peer Collaboration

**Collaborative discussion:** Agents interact in a non-hierarchical manner without predefined workflows, exchanging ideas, sharing information, and jointly exploring solutions. Specifically, agents operate on an equal level, allowing for free exchange of ideas and information to refine the reasoning processes of each other and reduce hallucinations in clinical contexts. There are no designated leader and workflow, promoting a more democratic and inclusive discussion. A common form of collaborative discussion is peer collaboration within a multidisciplinary team (MDT) structure, where agents role-play as experts from various disciplines. For instance, [31] proposed SeM-Agents, which includes different doctor roles and auxiliary roles to facilitate collaborative discussions. In this system, the doctor agent team is organized based on the specific situation of the patient, enabling multi-round discussions among multidisciplinary doctors. Once a consensus is reached, a summary agent reviews and presents the results. Another example of collaborative discussion in the human-computer interaction field is presented by [33]. They propose a multi-agent system for answering medical questions and predicting diagnoses, where a human physician collaborates with medical expert agents from diverse backgrounds. Together, they debate and generate diagnostic results, assisting the human physician in making comprehensive decisions.

**Clinical task-driven flow:** Agent communication aligns with the stages of a real-world clinical workflow, with information being passed and updated according to task progression. This flow typically includes stages such as triage, consultation, and diagnosis. At each stage, agents share relevant data, insights, and findings, enhancing decision-making and promoting efficient patient management. This structured approach streamlines communication and helps maintain focus on clinical objectives, ultimately improving patient outcomes. For example, [6] proposed MMedAgent-RL, which classifies agents into triage doctor, specialist doctor, and attending physician roles. The triage doctor agent performs initial departmental triage, routing patients to the appropriate department for

diagnosis. Specialist doctor agents discuss the patient's symptoms and provide diagnostic advice, while the attending physician agent makes the final diagnosis based on the discussions. This entire process follows a standard clinical task flow: triage → specialist consultation → attending doctor diagnosis.

**Problem-solving cycles:** Agent communication follows the phases of a problem-solving workflow, which can be either sequential or iterative. This cycle typically includes stages such as planning, execution, evaluation, and refinement. During each phase, agents collaborate by sharing information and insights to address challenges encountered in the diagnostic process. This collaboration enables them to assess the effectiveness of their actions and make necessary adjustments. For example, ClinicalAgent [4] has a goal of predicting clinical trial outcomes. The framework decompose the task into three sequential sub-tasks: task decomposition, subproblem solving, and reasoning. ClinicalAgent introduces a planning agent to decompose the clinical trial task into several sub-tasks and then recruit enrollment agent, safety agent, and efficacy agent to solve the sub-tasks from different perspectives. Finally, it provides a reasoning agent to make final decisions of the clinical trial outcomes. Rather than following a standard clinical task flow, ClinicalAgent follows a "divide-and-conquer" philosophy: it decomposes the whole process into three sub-stages and designs specific agents to solve each sub-task. Another example is HealthFlow [59], which provides a self-evolving iterative problem-solving workflow. A meta-agent takes task input and generate actionable plans. The generated plans are executed by executor agents and the execution results are processed by the evaluator agent to obtain feedback. A reflection agent takes the feedbacks and provides experience to the meta agent to refine the actionable plans. The whole workflow is task-oriented and self-evolving.

### 3. Discussion

In this section, we outline the main challenges and opportunities for advancing LLM-based multi-agent systems in medicine.

#### 3.1. Design and Evaluation of Agent Profiles

A critical step in building LLM-based multi-agent systems is the design of agent profiles. In medical contexts, this often involves assigning agents specific backgrounds, such as different specialties (e.g., radiology, pathology) or junior versus senior roles, reflecting the hierarchical and collaborative nature of real-world clinical practice. Despite its importance, research on how to design, evaluate, and measure these roles remains limited. For instance, current approaches, including role-playing with LLMs or fine-tuning on domain-specific corpora, aim to enhance the knowledge representation and role-playing capabilities of LLMs, yet it is unclear how faithfully they capture the subtleties of medical reasoning, inter-disciplinary dynamics, and decision-making authority. Evaluation should extend beyond task accuracy to assess role fidelity—for example, whether junior agents exercise appropriate caution or senior agents provide credible oversight. Role authenticity also remains challenging: agents may mimic the language of specialists without demonstrating genuine domain-consistent reasoning, producing superficially authoritative but potentially misleading outputs. Moreover, most systems assume static roles, whereas clinical responsibilities often shift dynamically depending on case complexity and team composition. Designing agent profiles that balance realism, adaptability, and reliability remains an open challenge.

#### 3.2. Self-Evolving Agentic Systems

Most existing multi-agent systems still rely on predefined architectures, static coordination strategies, and fixed agent-level configurations (e.g., knowledge bases, memory mechanisms, and reasoning pipelines), which inherently limit their ability to cope with change. The recent paradigm of self-evolving agents, however, highlights the possibility for agents to dynamically adapt and reorganize themselves. This capability is particularly important in medical contexts, where treatment decisions rarely remain static. As advances in medical technologies, the accumulation of clinical evidence, and the continuous revision of practice guidelines reshape therapeutic choices, agents

must be able to autonomously update their knowledge bases, reasoning pathways, and coordination mechanisms. A small number of studies have begun to refine agent coordination and communication through reinforcement learning [6], yet much more work is needed to explore continual adaptation mechanisms, long-term knowledge integration, and autonomous strategy evolution. Additionally, considering the stringent safety requirements in medical settings, an equally important challenge lies in ensuring that such updates are transparent, verifiable, and interpretable. Unlike other domains where autonomous adaptation may be tolerated with minimal oversight, in medicine every adjustment to an agent's knowledge or reasoning pipeline can directly influence patient outcomes. Thus, it is crucial that the evolution of agentic systems be accompanied by mechanisms that not only record and justify how new knowledge is incorporated, but also allow clinicians to audit, validate, and, when necessary, override these updates. Beyond interpretability at the decision level, transparency must also extend to the processes of coordination and knowledge integration across agents, so that the entire multi-agent system remains trustworthy and accountable.

### 3.3. Human Intervention

Much of the current research on LLM-based multi-agent systems in medical tasks has focused on fully automated approaches. However, such methods are inherently limited: they are vulnerable to hallucinations, cannot fully capture tacit clinical expertise, and risk reducing clinicians to passive validators. In high-stakes domains such as medicine, where safety and domain expertise are significant, human-in-the-loop mechanisms remain indispensable for ensuring clinical reliability, accountability, and trustworthiness [46,72,73]. A key limitation of current designs is that human involvement is often reduced to a final validation step, which treats clinicians as passive overseers rather than active collaborators. This raises several open questions: at what stages should human expertise be integrated; how much control should clinicians retain; and how can systems balance efficiency with safety? Without careful design, excessive reliance on physicians could increase workload, while too little involvement may erode trust. Future LLM-based multi-agent systems should therefore embrace human-AI collaboration as a design principle. Instead of restricting clinicians to a final validation step, systems should support interactive modes where experts can dynamically steer agent discussions, inject expertise knowledge, or arbitrate conflicts between divergent agent outputs. By advancing toward this participatory paradigm, multi-agent systems can move beyond static decision-support and evolve into partners in clinical reasoning.

### 3.4. Multimodal Integration

Clinical decision-making rarely depends on a single source of information. Instead, it synthesizes evidence from multiple modalities, such as imaging, genomic sequences, laboratory results, and electronic health records (EHRs). Some recent multi-agent medical systems have begun to incorporate multimodal inputs [6,22]. However, these systems handle these modalities as isolated inputs, each contributing to decision-making, without considering how they interact with one another. In practice, the interplay between modalities is far more complex. Multimodal information may exhibit different relationships: dominance (one modality outweighing others), complementarity (different modalities reinforcing each other), or conflict (modalities providing contradictory evidence) [74]. Thus, future research should move beyond simple aggregation and develop mechanisms to explicitly identify, reconcile, and leverage these relationships to support more robust and trustworthy multi-agent reasoning in clinical contexts.

### 3.5. Interaction Patterns

The way agents interact is central to the reliability of medical multi-agent systems. Existing work has largely explored two modes. Hierarchical coordination offers efficiency and accountability by assigning supervisory agents to recruit, delegate, and integrate, but it risks error propagation if oversight is flawed. Peer collaboration, in contrast, promotes balanced participation and robustness through mutual critique, yet often struggles with conflict resolution. Between these two extremes, a

widely adopted form is the multidisciplinary team (MDT), which mirrors real-world clinical consultations. MDTs can be understood as a structured instantiation of peer collaboration with elements of hierarchy: cases are presented, specialists contribute in turn, conflicting views are deliberated, and a chair physician synthesizes the outcome. However, current implementations rarely capture these procedural norms, often reducing MDT to unconstrained dialogue [5,7,42]. The key challenge, then, is not only to scale MDT-style collaboration but also to translate its well-established practices into agent workflows. Future work should investigate how to (1) design structured discussion protocols that mirror clinical MDT flow, (2) develop systematic mechanisms for resolving conflicting specialist outputs (e.g., weighted voting, confidence calibration, or arbitration), and (3) combine supervisory oversight with structured peer exchange so that accountability is preserved while ensuring that diverse expert reasoning is fully represented.

### 3.6. Medical Scenarios

Multi-agent systems have demonstrated growing application potential across a wide range of medical scenarios. Among these, diagnosis is the most extensively studied, ranging from general disease [5,7] to domain-specific settings such as glaucoma detection [37], cardiology [42], and rare disease diagnosis [34]. Beyond diagnosis, applications have also emerged in outpatient reception and triage [3,18], treatment planning and optimization [4,26,34], and clinical decision-making [45,54]. The objectives of these studies are diverse: some aim to reduce cognitive bias [33,60], others focus on mitigating hallucinations [25], improving explainability [43,54], or providing more personalized medical services [18]. Looking ahead, the advancement of multi-agent systems should expand to a broader range of medical tasks and scenarios that go beyond what clinicians can achieve. For instance, tasks such as efficacy and prognosis prediction in cancer immunotherapy, or the assessment of metastasis and recurrence risks are often difficult to address through clinical experience alone due to the high heterogeneity of tumors. These challenges are also suited to agent-based decision-making supported by models trained on large-scale medical datasets. Moreover, health education plays a critical role in enhancing public understanding of diseases and promoting early interventions. Multi-agent systems, equipped with role modeling and language style adaptation capabilities, can generate medical content that is both accessible and personalized based on the target audience. Compared to human physicians, agents are more effective at translating specialized medical content into comprehensible, public-facing language, thereby improving the dissemination efficiency of health information.

## 4. Conclusion

This survey systematically reviews the development of LLM-based multi-agent systems for medical problem-solving. We develop a medical-specific taxonomy along three dimensions: team composition, medical knowledge enhancement, and agent interactions, from our analysis of 50 papers. Despite recent progress, challenges remain in designing domain-specific agents and interactions. Future research should address these gaps, such as incorporating human-AI collaboration to ensure that human experts and multi-agent systems jointly address complex clinical tasks. We hope this work lays the groundwork for advancing reliable, impactful, and practically usable multi-agent systems in medicine.

## Appendix A. Appendix

### Appendix A.1. Acknowledgments of the Use of LLM

In this paper, we used ChatGPT to check grammar and improve wording. It did not change the original meaning of the text or introduce any new references or knowledge.

### Appendix A.2. Ethics Statement

Our survey does not involve human subjects, personal data, or sensitive information, and therefore does not raise direct ethical concerns. All analyzed papers are publicly available under appropriate licenses, and no private or identifiable information is included. We have carefully considered potential risks of bias, fairness, and misuse, and conclude that our work adheres to the ICLR Code of Ethics.

### Appendix A.3. Reproducibility Statement

We have taken concrete measures to ensure the reproducibility of our results. In particular, we provide the detailed encodings of all surveyed papers in Figure A1, along with comprehensive descriptions of the taxonomy construction process in the main text (Section 2). These materials collectively allow researchers to verify and extend our findings.



Figure A1. The coding of each paper under the proposed taxonomy.

## References

1. Chen, X.; Yi, H.; You, M.; Liu, W.; Wang, L.; Li, H.; Zhang, X.; Guo, Y.; Fan, L.; Chen, G.; et al. Enhancing diagnostic capability with multi-agents conversational large language models. *NPJ digital medicine* **2025**, *8*, 159.
2. Wang, Z.; Zhu, Y.; Zhao, H.; Zheng, X.; Sui, D.; Wang, T.; Tang, W.; Wang, Y.; Harrison, E.; Pan, C.; et al. Colacare: Enhancing electronic health record modeling through large language model-driven multi-agent collaboration. In Proceedings of the Proceedings of the ACM on Web Conference 2025, 2025, pp. 2250–2261.
3. Lu, M.; Ho, B.; Ren, D.; Wang, X. Triageagent: Towards better multi-agents collaborations for large language model-based clinical triage. In Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2024, 2024, pp. 5747–5764.
4. Yue, L.; Xing, S.; Chen, J.; Fu, T. Clinicalagent: Clinical trial multi-agent system with large language model-based reasoning. In Proceedings of the Proceedings of the 15th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics, 2024, pp. 1–10.
5. Kim, Y.; Park, C.; Jeong, H.; Chan, Y.S.; Xu, X.; McDuff, D.; Lee, H.; Ghassemi, M.; Breazeal, C.; Park, H.W. Mdagents: An adaptive collaboration of llms for medical decision-making. *Advances in Neural Information Processing Systems* **2024**, *37*, 79410–79452.
6. Xia, P.; Wang, J.; Peng, Y.; Zeng, K.; Wu, X.; Tang, X.; Zhu, H.; Li, Y.; Liu, S.; Lu, Y.; et al. MMedAgent-RL: Optimizing Multi-Agent Collaboration for Multimodal Medical Reasoning. *arXiv preprint arXiv:2506.00555* **2025**.
7. Wang, R.; Chen, Y.; Zhang, W.; Si, J.; Guan, H.; Peng, X.; Lu, W. MedConMA: A Confidence-Driven Multi-agent Framework for Medical Q&A. In Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining. Springer, 2025, pp. 421–433.
8. Li, X.; Wang, S.; Zeng, S.; Wu, Y.; Yang, Y. A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges. *Vicinagearth* **2024**, *1*, 9.
9. Guo, T.; Chen, X.; Wang, Y.; Chang, R.; Pei, S.; Chawla, N.V.; Wiest, O.; Zhang, X. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680* **2024**.
10. Yao, Z.; Yu, H. A survey on llm-based multi-agent ai hospital **2025**.
11. Alshehri, A.; Alshahrani, F.; Shah, H. A Precise Survey on Multi-agent in Medical Domains. *International Journal of Advanced Computer Science and Applications* **2023**, *14*. <https://doi.org/10.14569/IJACSA.2023.01406107>.
12. Wang, W.; Ma, Z.; Wang, Z.; Wu, C.; Ji, J.; Chen, W.; Li, X.; Yuan, Y. A survey of llm-based agents in medicine: How far are we from baymax? *arXiv preprint arXiv:2502.11211* **2025**.
13. Wang, Q.; Wang, T.; Tang, Z.; Li, Q.; Chen, N.; Liang, J.; He, B. MegaAgent: A large-scale autonomous LLM-based multi-agent system without predefined SOPs. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2025, 2025, pp. 4998–5036.
14. Zhuang, Y.; Jiang, W.; Zhang, J.; Yang, Z.; Zhou, J.T.; Zhang, C. Learning to Be A Doctor: Searching for Effective Medical Agent Architectures. *arXiv preprint arXiv:2504.11301* **2025**.
15. Almansoori, M.; Kumar, K.; Cholakkal, H. Self-Evolving Multi-Agent Simulations for Realistic Clinical Interactions. *arXiv preprint arXiv:2503.22678* **2025**.
16. Fan, Y.; Xue, K.; Li, Z.; Zhang, X.; Ruan, T. An LLM-based Framework for Biomedical Terminology Normalization in Social Media via Multi-Agent Collaboration. In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics, 2025, pp. 10712–10726.
17. Iapascurta, V.; Fiodorov, I.; Belii, A.; Bostan, V. Multi-Agent Approach for Sepsis Management. *Healthcare Informatics Research* **2025**, *31*, 209–214.
18. Bao, Z.; Liu, Q.; Guo, Y.; Ye, Z.; Shen, J.; Xie, S.; Peng, J.; Huang, X.; Wei, Z. Piors: Personalized intelligent outpatient reception based on large language model with multi-agents medical scenario simulation. *arXiv preprint arXiv:2411.13902* **2024**.
19. Zhao, W.; Wu, C.; Fan, Y.; Zhang, X.; Qiu, P.; Sun, Y.; Zhou, X.; Wang, Y.; Zhang, Y.; Yu, Y.; et al. An Agentic System for Rare Disease Diagnosis with Traceable Reasoning. *arXiv preprint arXiv:2506.20430* **2025**.
20. Wang, P.; Ye, S.; Naseem, U.; Kim, J. MRGAgents: A Multi-Agent Framework for Improved Medical Report Generation with Med-LVLMs. *arXiv preprint arXiv:2505.18530* **2025**.
21. Jia, Z.; Jia, M.; Duan, J.; Wang, J. DDO: Dual-Decision Optimization via Multi-Agent Collaboration for LLM-Based Medical Consultation. *arXiv preprint arXiv:2505.18630* **2025**.

22. Zhang, Y.; Bunting, K.V.; Champai, A.; Wang, X.; Lu, W.; Thorley, A.; Hothi, S.S.; Qiu, Z.; Kotecha, D.; Duan, J. CardAIC-Agents: A Multimodal Framework with Hierarchical Adaptation for Cardiac Care Support. *arXiv preprint arXiv:2508.13256* 2025.
23. Tang, X.; Zou, A.; Zhang, Z.; Li, Z.; Zhao, Y.; Zhang, X.; Cohan, A.; Gerstein, M. Medagents: Large language models as collaborators for zero-shot medical reasoning. *arXiv preprint arXiv:2311.10537* 2023.
24. Yan, W.; Liu, H.; Wu, T.; Chen, Q.; Wang, W.; Chai, H.; Wang, J.; Zhao, W.; Zhang, Y.; Zhang, R.; et al. Clinicallab: Aligning agents for multi-departmental clinical diagnostics in the real world. *arXiv preprint arXiv:2406.13890* 2024.
25. Low, C.H.; Wang, Z.; Zhang, T.; Zeng, Z.; Zhuo, Z.; Mazomenos, E.B.; Jin, Y. Surgraw: Multi-agent workflow with chain-of-thought reasoning for surgical intelligence. *arXiv preprint arXiv:2503.10265* 2025.
26. Chen, Z.; Peng, Z.; Liang, X.; Wang, C.; Liang, P.; Zeng, L.; Ju, M.; Yuan, Y. Map: Evaluation and multi-agent enhancement of large language models for inpatient pathways. *arXiv preprint arXiv:2503.13205* 2025.
27. Chen, Y.J.; Albarqawi, A.; Chen, C.S. Reinforcing clinical decision support through multi-agent systems and ethical ai governance. *arXiv preprint arXiv:2504.03699* 2025.
28. Ghezloo, F.; Seyfioglu, M.S.; Soraki, R.; Ikezogwo, W.O.; Li, B.; Vivekanandan, T.; Elmore, J.G.; Krishna, R.; Shapiro, L. Pathfinder: A multi-modal multi-agent system for medical diagnostic decision-making applied to histopathology. *arXiv preprint arXiv:2502.08916* 2025.
29. Zhou, Y.; Song, L.; Shen, J. MAM: Modular Multi-Agent Framework for Multi-Modal Medical Diagnosis via Role-Specialized Collaboration. *arXiv preprint arXiv:2506.19835* 2025.
30. Li, R.; Wang, X.; Berlowitz, D.; Mez, J.; Lin, H.; Yu, H. CARE-AD: a multi-agent large language model framework for Alzheimer's disease prediction using longitudinal clinical notes. *npj Digital Medicine* 2025, 8, 541.
31. Chen, K.; Qi, J.; Huo, J.; Tian, P.; Meng, F.; Yang, X.; Gao, Y. A Self-Evolving Framework for Multi-Agent Medical Consultation Based on Large Language Models. In Proceedings of the ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2025, pp. 1–5.
32. Mishra, P.P.; Arvan, M.; Zalake, M. TeamMedAgents: Enhancing Medical Decision-Making of LLMs Through Structured Teamwork. *arXiv preprint arXiv:2508.08115* 2025.
33. Li, H.; Cheng, X.; Zhang, X. Accurate Insights, Trustworthy Interactions: Designing a Collaborative AI-Human Multi-Agent System with Knowledge Graph for Diagnosis Prediction. In Proceedings of the Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, 2025, pp. 1–15.
34. Chen, X.; Jin, Y.; Mao, X.; Wang, L.; Zhang, S.; Chen, T. RareAgents: Advancing Rare Disease Care through LLM-Empowered Multi-disciplinary Team. *arXiv preprint arXiv:2412.12475* 2024.
35. Bao, L.; Peng, Z.; Zhou, X.; Cong, R.; Zhang, J.; Yuan, Y. Expertise-aware Multi-LLM Recruitment and Collaboration for Medical Decision-Making. *arXiv preprint arXiv:2508.13754* 2025.
36. Yang, D.; Wei, J.; Li, M.; Liu, J.; Liu, L.; Hu, M.; He, J.; Ju, Y.; Zhou, W.; Liu, Y.; et al. MedAide: Information Fusion and Anatomy of Medical Intents via LLM-based Agent Collaboration. *Information Fusion* 2025, p. 103743.
37. Liu, P.R.; Bansal, S.; Dinh, J.; Pawar, A.; Satishkumar, R.; Desai, S.; Gupta, N.; Wang, X.; Hu, S. Med-Chat: A Multi-Agent Framework for Multimodal Diagnosis with Large Language Models. *arXiv preprint arXiv:2506.07400* 2025.
38. Lyu, X.; Liang, Y.; Chen, W.; Ding, M.; Yang, J.; Huang, G.; Zhang, D.; He, X.; Shen, L. WSI-Agents: A Collaborative Multi-Agent System for Multi-Modal Whole Slide Image Analysis. *arXiv preprint arXiv:2507.14680* 2025.
39. Zhao, Y.; Wang, H.; Zheng, Y.; Wu, X. A Layered Debating Multi-Agent System for Similar Disease Diagnosis. In Proceedings of the Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers), 2025, pp. 539–549.
40. Solovev, G.V.; Zhidkovskaya, A.B.; Orlova, A.; Vepreva, A.; Ilya, T.; Golovinskii, R.; Gubina, N.; Chistakov, D.; Aliev, T.A.; Poddiakov, I.; et al. Towards LLM-driven multi-agent pipeline for drug discovery: neurodegenerative diseases case study. In Proceedings of the 2nd AI4Research Workshop: Towards a Knowledge-grounded Scientific Research Lifecycle, 2024.
41. Zuo, K.; Zhong, Z.; Huang, P.; Tang, S.; Chen, Y.; Jiang, Y. HEAL-KGGen: A Hierarchical Multi-Agent LLM Framework with Knowledge Graph Enhancement for Genetic Biomarker-Based Medical Diagnosis. *bioRxiv* 2025, pp. 2025–06.

42. Zhang, W.; Qiao, M.; Zang, C.; Niederer, S.; Matthews, P.M.; Bai, W.; Kainz, B. Multi-Agent Reasoning for Cardiovascular Imaging Phenotype Analysis. *arXiv preprint arXiv:2507.03460* **2025**.
43. Hong, S.; Xiao, L.; Zhang, X.; Chen, J. ArgMed-Agents: explainable clinical decision reasoning with LLM discussion via argumentation schemes. In Proceedings of the 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). IEEE, 2024, pp. 5486–5493.
44. Liu, S.; Lu, Y.; Chen, S.; Hu, X.; Zhao, J.; Lu, Y.; Zhao, Y. Drugagent: Automating ai-aided drug discovery programming through llm multi-agent collaboration. *arXiv preprint arXiv:2411.15692* **2024**.
45. Ke, Y.H.; Yang, R.; Lie, S.A.; Lim, T.X.Y.; Abdullah, H.R.; Ting, D.S.W.; Liu, N. Enhancing diagnostic accuracy through multi-agent conversations: using large language models to mitigate cognitive bias. *arXiv preprint arXiv:2401.14589* **2024**.
46. Xu, H.; Yi, S.; Lim, T.; Xu, J.; Well, A.; Mery, C.; Zhang, A.; Zhang, Y.; Ji, H.; Pingali, K.; et al. Tama: A human-ai collaborative thematic analysis framework using multi-agent llms for clinical interviews. *arXiv preprint arXiv:2503.20666* **2025**.
47. Low, C.H.; Zhuo, Z.; Wang, Z.; Xu, J.; Liu, H.; Sirajudeen, N.; Boal, M.; Edwards, P.J.; Stoyanov, D.; Francis, N.; et al. CARES: Collaborative Agentic Reasoning for Error Detection in Surgery. *arXiv preprint arXiv:2508.08764* **2025**.
48. Zhao, H.; Zhu, Y.; Wang, Z.; Wang, Y.; Gao, J.; Ma, L. ConfAgents: A Conformal-Guided Multi-Agent Framework for Cost-Efficient Medical Diagnosis. *arXiv preprint arXiv:2508.04915* **2025**.
49. Wu, H.; Zhu, Y.; Wang, Z.; Zheng, X.; Wang, L.; Tang, W.; Wang, Y.; Pan, C.; Harrison, E.M.; Gao, J.; et al. EHRFlow: A Large Language Model-Driven Iterative Multi-Agent Electronic Health Record Data Analysis Workflow. In Proceedings of the KDD'24 Workshop: Artificial Intelligence and Data Science for Healthcare: Bridging Data-Centric AI and People-Centric Healthcare, 2024.
50. Liang, C.; Ma, Z.; Wang, W.; Ding, M.; Cao, Z.; Chen, M. MCM: A Multi-Agent Collaborative Multimodal Framework For Traditional Chinese Medicine Diagnosis. In Proceedings of the 2025 IEEE International Conference on Image Processing (ICIP). IEEE, 2025, pp. 1438–1443.
51. Zhou, Y.; Zhang, P.; Song, M.; Zheng, A.; Lu, Y.; Liu, Z.; Chen, Y.; Xi, Z. Zodiac: A cardiologist-level llm framework for multi-agent diagnostics. *arXiv preprint arXiv:2410.02026* **2024**.
52. Zhou, X.; Ren, Y.; Zhao, Q.; Huang, D.; Wang, X.; Zhao, T.; Zhu, Z.; He, W.; Li, S.; Xu, Y.; et al. An LLM-Driven Multi-Agent Debate System for Mendelian Diseases. *arXiv preprint arXiv:2504.07881* **2025**.
53. Xie, Y.; Cui, H.; Zhang, Z.; Lu, J.; Shu, K.; Nahab, F.; Hu, X.; Yang, C. KERAP: A Knowledge-Enhanced Reasoning Approach for Accurate Zero-shot Diagnosis Prediction Using Multi-agent LLMs. *arXiv preprint arXiv:2507.02773* **2025**.
54. Liu, Z.; Xiao, L.; Zhu, R.; Yang, H.; He, M. MedGen: An Explainable Multi-Agent Architecture for Clinical Decision Support through Multisource Knowledge Fusion. In Proceedings of the 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM). IEEE, 2024, pp. 6474–6481.
55. Shi, H.; Zhang, J.; Zhang, K. Enhancing Clinical Trial Patient Matching through Knowledge Augmentation and Reasoning with Multi-Agents. *arXiv preprint arXiv:2411.14637* **2024**.
56. Chen, K.; Zhen, T.; Wang, H.; Liu, K.; Li, X.; Huo, J.; Yang, T.; Xu, J.; Dong, W.; Gao, Y. MedSentry: Understanding and Mitigating Safety Risks in Medical LLM Multi-Agent Systems. *arXiv preprint arXiv:2505.20824* **2025**.
57. Mahajan, B.; Ji, K. MediLink: A Multi-Agent Conversational Pipeline for Evidence-Grounded Symptom Diagnosis **2025**.
58. Chen, C.; Weishaupt, L.L.; Williamson, D.F.; Chen, R.J.; Ding, T.; Chen, B.; Vaidya, A.; Le, L.P.; Jaume, G.; Lu, M.Y.; et al. Evidence-based diagnostic reasoning with multi-agent copilot for human pathology. *arXiv preprint arXiv:2506.20964* **2025**.
59. Zhu, Y.; Qi, Y.; Wang, Z.; Gu, L.; Sui, D.; Hu, H.; Zhang, X.; He, Z.; Ma, L.; Yu, L. HealthFlow: A Self-Evolving AI Agent with Meta Planning for Autonomous Healthcare Research. *arXiv preprint arXiv:2508.02621* **2025**.
60. Ke, Y.; Yang, R.; Lie, S.A.; Lim, T.X.Y.; Ning, Y.; Li, I.; Abdullah, H.R.; Ting, D.S.W.; Liu, N. Mitigating cognitive biases in clinical decision-making through multi-agent conversations using large language models: simulation study. *Journal of Medical Internet Research* **2024**, *26*, e59439.
61. Liang, Y.; Lyu, X.; Chen, W.; Ding, M.; Zhang, J.; He, X.; Wu, S.; Xing, X.; Yang, S.; Wang, X.; et al. WSI-LLaVA: A multimodal large language model for whole slide image. *arXiv preprint arXiv:2412.02141* **2024**.
62. Seyfioglu, M.S.; Ikezogwo, W.O.; Ghezloo, F.; Krishna, R.; Shapiro, L. Quilt-llava: Visual instruction tuning by extracting localized narratives from open-source histopathology videos. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 13183–13192.

63. Sellergren, A.; Kazemzadeh, S.; Jaroensri, T.; Kiraly, A.; Traverse, M.; Kohlberger, T.; Xu, S.; Jamil, F.; Hughes, C.; Lau, C.; et al. Medgemma technical report. *arXiv preprint arXiv:2507.05201* **2025**.
64. Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* **2024**.
65. Luo, Y.; Zhang, J.; Fan, S.; Yang, K.; Wu, Y.; Qiao, M.; Nie, Z. Biomedgpt: Open multimodal generative pre-trained transformer for biomedicine. *arXiv preprint arXiv:2308.09442* **2023**.
66. Demner-Fushman, D.; Kohli, M.D.; Rosenman, M.B.; Shooshan, S.E.; Rodriguez, L.; Antani, S.; Thoma, G.R.; McDonald, C.J. Preparing a collection of radiology examinations for distribution and retrieval. *Journal of the American Medical Informatics Association* **2015**, *23*, 304–310.
67. Johnson, A.E.; Pollard, T.J.; Greenbaum, N.R.; Lungren, M.P.; Deng, C.y.; Peng, Y.; Lu, Z.; Mark, R.G.; Berkowitz, S.J.; Horng, S. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs. *arXiv preprint arXiv:1901.07042* **2019**.
68. Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y.J.; Madotto, A.; Fung, P. Survey of hallucination in natural language generation. *ACM computing surveys* **2023**, *55*, 1–38.
69. Smedley, D.; Jacobsen, J.O.; Jäger, M.; Köhler, S.; Holtgrewe, M.; Schubach, M.; Siragusa, E.; Zemojtel, T.; Buske, O.J.; Washington, N.L.; et al. Next-generation diagnostics and disease-gene discovery with the Exomiser. *Nature protocols* **2015**, *10*, 2004–2015.
70. Mao, X.; Huang, Y.; Jin, Y.; Wang, L.; Chen, X.; Liu, H.; Yang, X.; Xu, H.; Luan, X.; Xiao, Y.; et al. A phenotype-based AI pipeline outperforms human experts in differentially diagnosing rare diseases using EHRs. *npj Digital Medicine* **2025**, *8*, 68.
71. Shen, C.; Zhang, W.; Li, K.; Huang, E.; Bi, H.; Fan, A.; Shen, Y.; Dong, H.; Zhang, J.; Shao, Y.; et al. FEAT: A Multi-Agent Forensic AI System with Domain-Adapted Large Language Model for Automated Cause-of-Death Analysis. *arXiv preprint arXiv:2508.07950* **2025**.
72. Sheng, R.; Shi, C.; Lotfi, S.; Liu, S.; Perer, A.; Qu, H.; Cheng, F. Design Patterns of Human-AI Interfaces in Healthcare. *arXiv preprint arXiv:2507.12721* **2025**.
73. Sheng, R.; Wang, X.; Wang, J.; Jin, X.; Sheng, Z.; Xu, Z.; Rajendran, S.; Qu, H.; Wang, F. TrialCompass: Visual Analytics for Enhancing the Eligibility Criteria Design of Clinical Trials. *arXiv preprint arXiv:2507.12298* **2025**.
74. Wang, X.; He, J.; Jin, Z.; Yang, M.; Wang, Y.; Qu, H. M2lens: Visualizing and explaining multimodal models for sentiment analysis. *IEEE Transactions on Visualization and Computer Graphics* **2021**, *28*, 802–812.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.