

LungNoduleAgent: A Collaborative Multi-Agent System for Precision Diagnosis of Lung Nodules

Cheng Yang¹, Hui Jin¹, Xinlei Yu², Zhipeng Wang¹, Yaoqun Liu³, Fenglei Fan⁴,
Dajiang Lei⁵, Gangyong Jia¹, Changmiao Wang⁶, Ruiquan Ge^{1,*}

¹ Hangzhou Dianzi University, Hangzhou, China

² National University of Singapore, Singapore

³ The Chinese University of Hong Kong, Shenzhen, Shenzhen, China

⁴ City University of Hong Kong, Hong Kong, China

⁵ Chongqing University of Posts and Telecommunications, Chongqing, China

⁶ Shenzhen Research Institute of Big Data, Shenzhen, China

Abstract

Diagnosing lung cancer typically involves physicians identifying lung nodules in Computed tomography (CT) scans and generating diagnostic reports based on their morphological features and medical expertise. Although advancements have been made in using multimodal large language models for analyzing lung CT scans, challenges remain in accurately describing nodule morphology and incorporating medical expertise. These limitations affect the reliability and effectiveness of these models in clinical settings. Collaborative multi-agent systems offer a promising strategy for achieving a balance between generality and precision in medical applications, yet their potential in pathology has not been thoroughly explored. To bridge these gaps, we introduce LungNoduleAgent, an innovative collaborative multi-agent system specifically designed for analyzing lung CT scans. LungNoduleAgent streamlines the diagnostic process into sequential components, improving precision in describing nodules and grading malignancy through three primary modules. The first module, the Nodule Spotter, coordinates clinical detection models to accurately identify nodules. The second module, the Radiologist, integrates localized image description techniques to produce comprehensive CT reports. Finally, the Doctor Agent System performs malignancy reasoning by using images and CT reports, supported by a pathology knowledge base and a multi-agent system framework. Extensive testing on two private datasets and the public LIDC-IDRI dataset indicates that LungNoduleAgent surpasses mainstream vision-language models, agent systems, and advanced expert models such as GPT-4o, Claude 3.7 Sonnet, LLaMA-3.2 Vision, Qwen2.5-VL, Med-R1, MedGemma, MedAgent-Pro, MedAgents, MDAgent and LLaVA-Med. These results highlight the importance of region-level semantic alignment and multi-agent collaboration in diagnosing nodules. LungNoduleAgent stands out as a promising foundational tool for supporting clinical analyses of lung nodules.

Code — <https://github.com/ImYangC7/LungNoduleAgent>

*Corresponding author: Ruiquan Ge, gespring@hdu.edu.cn.
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Introduction

Lung cancer remains a leading cause of cancer-related deaths globally, with early detection and accurate diagnosis being vital for improving patient outcomes (Tammemagi et al. 2019). Computed tomography (CT) scans are crucial for identifying lung nodules, which serve as early indicators of malignancy (Swanson et al. 2023). Traditionally, radiologists examine these scans by evaluating nodule morphology and applying their medical expertise to produce diagnostic reports (Osarogiagbon et al. 2023). However, this process necessitates radiologists to manually examine each subsequent CT image (Hammer et al. 2019; Lee et al. 2024), which is time-consuming and susceptible to interobserver variability (Driessen et al. 2025).

In the realm of lung cancer analysis using CT images, deep learning advancements have greatly improved the screening and diagnosis of nodules (Swanson et al. 2023). These technological strides primarily focus on three key tasks: cancer classification (Nakach, Idri, and Goceri 2024; Raza et al. 2023; Ji et al. 2023), grading (Fan et al. 2025; Shen et al. 2025; Martinez-Murcia et al. 2021; Wang et al. 2018), and nodule detection (Cao et al. 2020; UrRehman et al. 2024). Despite these achievements, deep learning models still face significant challenges (Ladbury et al. 2023). One major issue is the interpretability of model outputs. While these models often achieve high performance metrics, their reasoning processes are not easily understood, which impedes their acceptance in clinical settings. Unlike traditional clinical reasoning, which follows transparent diagnostic procedures that physicians can explain to patients, deep learning models typically offer conclusions without revealing the specific features or patterns that underpin their predictions. Consequently, this lack of transparency poses a barrier to their integration into clinical practice. Another challenge is the dependency of deep learning models on diverse datasets. When confronted with new, unseen data, these models may fail to generalize effectively. This limitation highlights the necessity for models that are adaptable to various scenarios. Furthermore, many current methods are designed for specific tasks, reducing their flexibil-

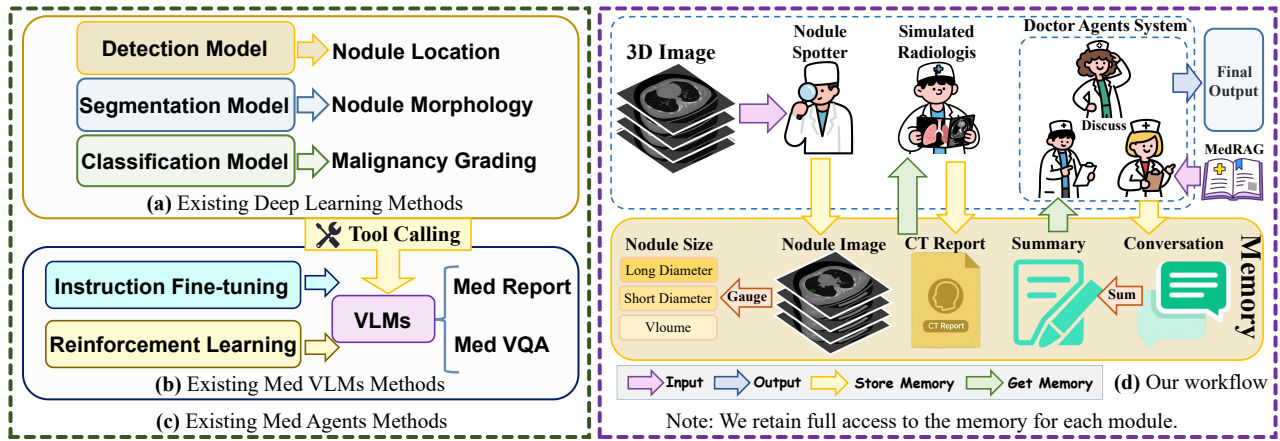


Figure 1: Workflow Comparison of Traditional Methods vs. the Proposed LungNoduleAgent Framework.

ity. This situation underscores the urgent need for comprehensive models capable of addressing multiple pathological tasks while also providing enhanced interpretability.

Recently, newly released general Vision-Language Models (VLMs) (Zhang et al. 2024) such as LLaVA (Liu et al. 2023a), GPT-4o (Achiam et al. 2023), Qwen2.5VL (Bai et al. 2025), InternVL (Zhu et al. 2025) and Claude 3.7 Sonnet (Anthropic 2024) have shown impressive interpretation and generalization abilities. However, these models often fall short in specialized medical contexts due to insufficient domain-specific training and a lack of domain prior knowledge, which prevents them from meeting the demands of professional medical applications. Medical VLMs (Thirunavukarasu et al. 2023), including MedR1 (Lai et al. 2025), LLaVA-Med (Li et al. 2023), PMC-VQA (Zhang et al. 2023), and MedGemma (Selligren et al. 2025), which are based on supervised fine-tuning and reinforcement learning, have improved their multimodal reasoning and generalization capabilities in the medical field. Nevertheless, their limited fine-grained visual perception hinders quantitative analysis. These models often rely on the internal knowledge of VLMs for judgments, while modern medical practice emphasizes evidence-based diagnosis, which requires structured reasoning and clinical evidence. Simultaneously, agentic systems (Liu et al. 2023b; Yang et al. 2025; Yu et al. 2025) that use collaborative multi-agent systems (Liang et al. 2023; Du et al. 2023; Chan et al. 2023) for clinical reasoning or integrate external DL tools (Wang et al. 2025; Li et al. 2024a; Fallahpour et al. 2025) to extend the capabilities of VLMs offer a promising solution. However, their application in lung cancer diagnosis is not yet mature, lacking fine-grained analysis of lung nodules and sufficient pathology-specific knowledge. This leads to an accuracy of only 40-50% on lung cancer-specific tasks, while achieving 75-80% (Tang et al. 2023; Kim et al. 2024; Wang et al. 2024; Liu et al. 2024) on general medical tasks. This underscores the need for a more comprehensive and collaborative diagnostic framework in the field of nodule analysis (Qiu et al. 2025).

To address these challenges, we introduce LungNod-

uleAgent, the first collaborative multi-agent system for lung nodule analysis, designed to enhance the accuracy and reliability of lung CT scan analysis. As shown in **Figure 1**, our contributions are as follows:

- The system mimics the clinical workflow by decomposing the diagnostic process into sequential stages, each handled by a specialized component: “Nodule Spotter” focuses on precise nodule detection, “Simulated Radiologist” generates detailed CT reports, and “Doctor Agent System (DAS)” assesses malignancy by leveraging expert knowledge and multi-agent discussions.
- We employ a Focal Prompting Mechanism to enhance the model’s fine-grained visual perception, enabling the description of nodule morphological features and the dynamic evaluation of nodule characteristics across slices.
- By storing pathological information in a shared Memory for the agent system, we utilize collaborative multi-agent discussions and introduce medical prior knowledge to achieve evidence-based quantitative analysis and fact-based clinical reasoning.

Comprehensive evaluations on two private and one public dataset show our system is significantly more effective than existing VLMs and medical agentic systems in nodule report generation and malignancy assessment.

Methodology

Overview Architecture

The workflow of LungNoduleAgent, illustrated in **Figure 2**, processes lung CT volume V , through three modules: Nodule Spotter, Simulated Radiologist, Doctor Agent System.

The Nodule Spotter begins by dividing the CT volume into slices I , which are then processed through a Mixture of Experts architecture to produce initial masks m . These masks go through Mask Clustering, which refines them into candidate masks M_g and excludes outliers. A Judging Panel uses VLMs to validate M_g , determining the final mask M .

In the Simulated Radiologist module, the image slice I and the final mask M are encoded using a Focal Prompting

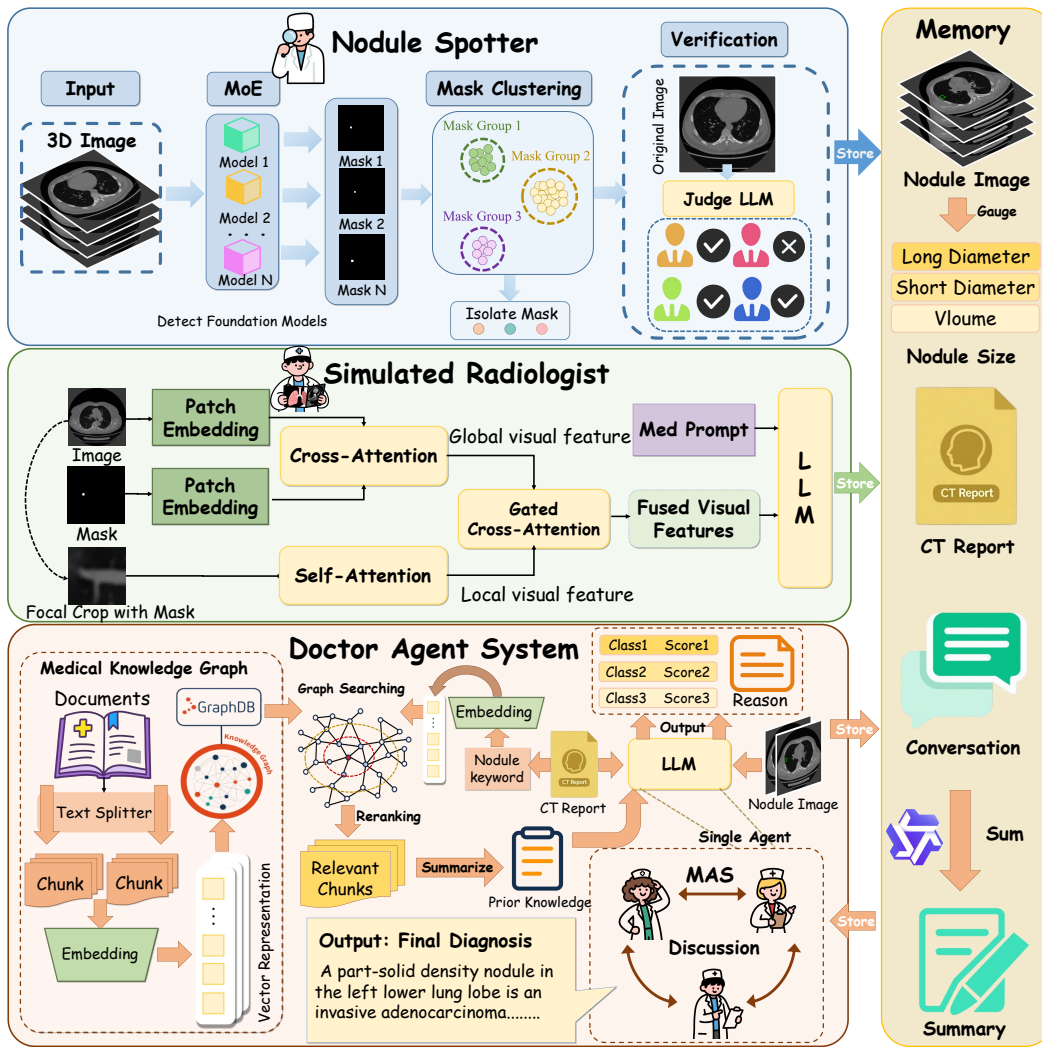


Figure 2: Overview of LungNoduleAgent for multi-modal lung nodule analysis with a nodule spotter for lung nodule detection, simulated radiologist for localized CT report generation, and Doctor Agent System for malignancy grading.

Mechanism that highlights lung nodule regions while retaining surrounding context. A MedPrompt further refines the generation of CT reports, ensuring focus on annotated areas and the integration of visual and textual information to produce localized CT Reports. Both the Nodule Image and the CT Report are stored in the Memory module for future access and interaction.

The Doctor Agent System employs each Medical Agent as a VLM enhanced with a medical knowledge graph. This integration of domain-specific medical knowledge enriches the VLMs with relevant pathological information, aiding in constructing reasoning chains and providing diagnostic explanations. Initially, each agent, with unique domain expertise, analyzes the Nodule Image and CT Report to generate its diagnosis and supporting rationale, stored in the Memory’s Conversation module to generate. A designated VLM then compiles these initial outputs into a summary, which is used for a second round of discussion among the agents.

This iterative process continues until consensus is reached, culminating in the final diagnosis FD .

Nodule Spotter

The Nodule Spotter module is essential in the diagnostic process, as it identifies and pinpoints lung nodule regions. This capability allows radiologists to create descriptions that are tailored to these localized findings.

Mixture of Experts. This submodule processes 3D images via slice-based division for analysis. It adopts a Mixture of Experts (MoE) framework, incorporating multiple specialized foundational detection models, each proficient in identifying specific nodule types or features. These models operate in parallel, analyzing individual lung CT slices independently, with each expert generating a m to denote predicted nodule locations. Such parallel processing leverages the strengths of diverse models, enhancing the robustness and comprehensiveness of initial nodule detection through

the integration of their outputs.

Mask Clustering. This module clusters initial masks based on their spatial overlap to obtain a more stable and accurate final mask representation. Masks with substantial overlap are grouped into the same cluster, while those with insufficient similarity to any group are treated as outliers and removed.

To quantify mask similarity, we define the distance between two masks m_i and m_j as

$$d(m_i, m_j) = 1 - \text{IoU}(m_i, m_j), \quad (1)$$

where the Intersection over Union is

$$\text{IoU}(m_i, m_j) = \frac{|m_i \cap m_j|}{|m_i \cup m_j|}, \quad (2)$$

and $|\cdot|$ denotes the number of pixels. Distances therefore lie in $[0, 1]$, with smaller values indicating stronger spatial overlap. This metric forms the basis for mask grouping.

We adopt DBSCAN for clustering, parameterized by a neighborhood radius $\epsilon \in [0, 1]$ —equivalent to an IoU threshold $\tau = 1 - \epsilon$ —and a minimum sample count MinPts . The ϵ -neighborhood of a mask is defined as

$$N_\epsilon(m_i) = \{m_j \mid d(m_i, m_j) \leq \epsilon\}. \quad (3)$$

DBSCAN marks all masks as unvisited initially. For each mask m_i , if $|N_\epsilon(m_i)| \geq \text{MinPts}$, a new cluster is created and expanded by recursively adding all masks reachable through connected ϵ -neighborhoods. Masks that do not meet this density requirement are labeled noise. The algorithm outputs clusters $C = \{C_1, \dots, C_K\}$ and noise set N .

For each cluster $C_k = \{m_1, \dots, m_N\}$, we compute an averaged mask \bar{m}_k to aggregate the spatial consensus among its members. The averaged mask is then binarized with a threshold of 0.5, producing the final refined mask used for subsequent nodule analysis.

Judging Panel. The Judging Panel validation submodule is the crucial final step in distinguishing true nodule candidates from false positives, utilizing a consensus approach among VLMs. For each nodule candidate, represented by the mask group M_g and the original image slice I , N_{VLM} independent VLMs simultaneously assess the visual alignment between the mask and anatomical features. After undergoing evaluation by the Judging Panel, the final mask M is determined.

Each VLM, denoted as \mathcal{V}_j , performs its assessment through two outputs: (1) a binary decision, $\text{Sign}(\mathcal{V}_j)$, where $+1$ indicates approval as a valid nodule, and -1 indicates rejection, as defined by:

$$\text{Sign}(\mathcal{V}_j) = \begin{cases} +1, & \text{approval,} \\ -1, & \text{rejection,} \end{cases} \quad (4)$$

and (2) a confidence score C_j that quantifies the certainty of the alignment.

The system consolidates these evaluations through weighted scoring:

$$\text{Score}(M_g) = \sum_{j=1}^{N_{VLM}} (\text{Sign}(\mathcal{V}_j) \times C_j), \quad (5)$$

where a $\text{Score}(M_g) > 0$ confirms the validity of the candidate through majority consensus.

This design, combining discrete decisions with continuous confidence metrics, simulates a rigorous peer-review process. It mitigates individual model biases while establishing interpretable thresholds suitable for clinical applications. Nodule candidates deemed invalid, with $\text{TotalScore} \leq 0$, are automatically filtered out as false positives, ensuring that only anatomically accurate detections are considered for further analysis.

Simulated Radiologist

Focal Prompting Mechanism. Simulated Radiologist utilizes a focal prompting technique inspired by the Describe Anything Model (Zhang et al. 2021) to encode areas with lung nodules while preserving the surrounding context for thorough analysis. This technique involves focal cropping of both the image and its corresponding mask, ensuring that nearby areas are maintained for local context.

Both the full image and the focal crop are processed by a local visual backbone, embedding the image and binary mask in a spatially aligned manner. To deepen the understanding of the focal crop, global context from the full image is integrated through gated cross-attention. Through sequence concatenation, the dynamic changes of nodules across slices are captured to enhance the overall perception of the nodules (Mao et al. 2025). This approach enhances the detailed comprehension and analysis of the nodule regions within the broader anatomical context.

Region-specific Prompting. We employ MedPrompt to guide the effective generation of CT reports. Unlike general prompts, MedPrompt ensures that the model focuses exclusively on annotated areas, utilizing anatomically accurate terminology and maintaining a professional clinical report format while avoiding speculative or irrelevant content. It incorporates clearly defined output formats and customized objectives to reduce errors and inaccurate descriptions. The fused feature, combined with the prompt, is then processed by VLM:

$$\mathcal{O}_{vlm} = \text{VLM}(\text{MedPrompt}, \Theta_{\text{volume}}). \quad (6)$$

This process ensures the production of accurate and clinically relevant CT reports \mathcal{O}_{vlm} .

Doctor Agent System

Medical Graph RAG. To address the limited pathological knowledge of existing medical agents, we have adopted a medical graph construction and retrieval strategy inspired by GraphRag (Edge et al. 2024), enabling these agents to access relevant pathological information. This resource aids in constructing reasoning chains and providing diagnostic justifications based on CT reports and medical knowledge. While many medical large language models utilize Retrieval-Augmented Generation (RAG) (Arslan et al. 2024) to obtain external knowledge, they often encounter difficulties with broad queries, such as identifying the primary characteristics of malignant lung nodules, because

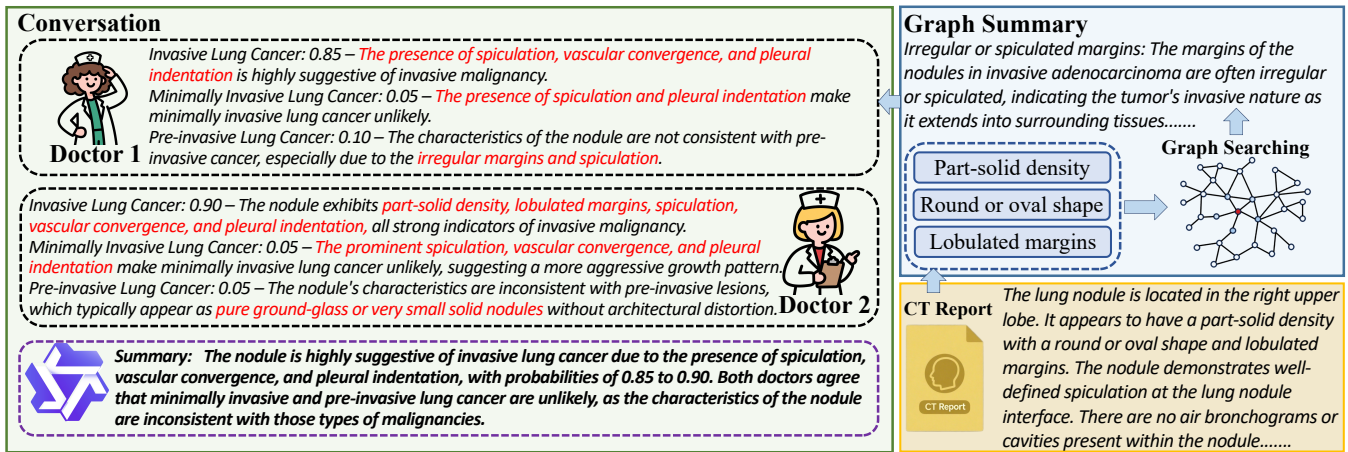


Figure 3: Visualization of the DAS’s Internal Collaboration. This figure illustrates how multiple agents leverage a medical knowledge graph and a collaborative conversational mechanism to infer and arrive at a final diagnosis for lung nodules based on the CT Report and Nodule Image.

these questions necessitate query-focused summarization rather than straightforward retrieval.

To address this issue, we propose a Graph Searching framework that compiles data from authoritative pathological websites and literature to create a comprehensive knowledge base. Specifically, the knowledge graph \mathcal{G} is constructed from documents \mathcal{D} :

$$\mathcal{G} = \text{GraphConstruct}(\mathcal{D}). \quad (7)$$

Community-level summaries \mathcal{S} are then derived from the graph. When presented with a user query Q , the system uses a language model to generate the final answer \mathcal{A} :

$$\mathcal{A} = \text{VLM}(\mathcal{S}, Q, \mathcal{N}), \quad (8)$$

where \mathcal{N} represents the nodule image, defined as the image I covered with mask M .

Through this graph-based knowledge infusion pipeline ($\mathcal{D} \rightarrow \mathcal{G} \rightarrow \mathcal{S} \rightarrow \mathcal{A}$), the large language model evolves into a reasoning agent with specialized medical expertise. This transformation enhances its ability to deliver accurate and detailed diagnostic insights.

Multi-Agent System. Our multi-agent system, illustrated in **Figure 3**, employs a decentralized and collaborative approach akin to a roundtable discussion. This methodology is designed to enhance diagnoses, verify information, and ensure the precision of reasoning. At the outset, K reasoning agents independently assess the report and image I , each generating their initial outputs:

$$O_i^{(1)} = \text{Agent}_i(I, \text{Report}), \quad i = 1, \dots, K, \quad (9)$$

where i denotes the i th agent. If discrepancies occur, agents refine their conclusions by considering the perspectives of other agents:

$$O_i^{(t)} = \text{Revise}(O_i^{(t-1)}, \{O_j^{(t-1)}\}_{j \neq i}). \quad (10)$$

Subsequently, a summarization agent consolidates these outputs into a cohesive summary. This iterative process continues until a consensus is achieved:

$$\mathcal{FD} = \text{Consensus} \left(\text{Summarizer} \left(\{O_i^{(t^*)}\}_{i=1}^K \right) \right), \quad (11)$$

where \mathcal{FD} represents the final diagnosis.

Through this multi-round communication and refinement, the system delivers reliable diagnostic results grounded in collective agreement among the agents.

Memory

The Memory Module serves as the system’s core storage component, designed for managing critical information. It stores nodule images, while measuring and storing nodule size as evidence. It also holds CT Reports, which are initial interpretations of visual image features. During the DAS discussion process, the Memory Module stores detailed Conversations and Summaries from multi-agent interactions. This layered and real-time memory mechanism facilitates a more intuitive understanding of lung nodules by the agents and reduces information processing complexity.

Experiments

Datasets and Implementation

Our experiments utilized two private datasets comprised of 1,616 and 386 axial lung CT slices (512×512 pixels). These datasets are carefully annotated to aid lung nodule detection, featuring bounding box masks, comprehensive morphological descriptions including lobar location, density, shape, margin, and indications of cavitation or vacuolation, and classifications of malignancy, categorized as pre-invasive, minimally invasive, or invasive adenocarcinoma. In addition to these private sources, we employed the public dataset LIDC-IDRI (Samuel 2011), including 1,018 lung CT scans with identical resolution. This dataset provides detailed lung nodule segmentation masks and semantic attributes such as

size, margin clarity, density, and spiculation. Furthermore, it offers malignancy ratings ranging from 1 to 5, classifying scores above 3 as malignant and those below as benign.

Attribute Question Construction

To address the challenge of lacking precise ground truth for localized CT reports, we adopt an attribute-level verification approach inspired by the MedDLC-score (Xiao et al. 2025). For each nodule, we develop a series of clinically relevant yes/no questions based on dataset annotations to evaluate if the generated descriptions accurately capture key morphological traits such as shape, margin, and size. These questions are split into two categories: positive questions (Pos QA) verify the presence of specific nodule characteristics, while negative questions (Neg QA) detect irrelevant or fabricated details. The average correctness rate of these questions yields the LungDLC-score, a reliable metric for assessing the accuracy with which reports capture essential clinical features and morphology.

Comparison Methods and Evaluation Metrics

We assessed our proposed method against a selection of current models, including general-purpose VLMs like GPT-4o (Achiam et al. 2023), Claude 3.7 Sonnet (Anthropic 2024), InternVL (Zhu et al. 2025), LLaVA3.2 (Liu et al. 2023a), and Qwen2.5-VL (Bai et al. 2025). Additionally, we evaluated medical agents such as Med-R1 (Lai et al. 2025), MedGamma (Sellergren et al. 2025), MedAgent-Pro (Wang et al. 2025), MDAgents (Kim et al. 2024), MedAgents (Tang et al. 2023), and LLaVA-Med (Li et al. 2023). These models epitomize the latest in visual-language pre-training and domain-specific tuning advancements. To ensure consistency, all baseline models were given full-slice lung CT images and the same MedPrompt.

We evaluate CT report generation using a two-pronged approach. First, we utilize “LLM-as-a-judge” framework (Li et al. 2024b), where GPT-4o rates generated reports for fluency, relevance, consistency, and clinical rationality. Each aspect is scored independently, and their average forms the LLM-score. Second, to reduce reliance on ground truth, we introduce the LungDLC-score, which checks whether descriptions capture nodule features and avoid irrelevant or fabricated content. For malignancy grading, we report Accuracy (Acc) and F1-score.

Results and Discussions

Results on localized CT report generation task. As demonstrated in **Table 1**, LungNoduleAgent exhibits state-of-the-art performance across all three benchmark datasets, with significant improvements on the Qwen and LLaVA base models. Compared with the highest score among other methods, on the PrivateA dataset, it achieves a LungDLC-score of 81.9, reflecting an increase of 6.3. For the PrivateB dataset, the model attains a score of 80.3, marking an improvement of 4.1. This performance advantage is also evident on the public LIDC-IDRI benchmark, where the model scores 83.5, representing 8.3 increasing.

Notably, the agent demonstrates balanced capabilities in both Positive QA and Negative QA. While GPT-4o shows

marginally higher LLM-scores due to evaluator bias (Li et al. 2025), our model’s consistent superiority in objective metrics validates its clinical utility.

The result of the malignancy grading task. LungNoduleAgent achieved state-of-the-art performance across all evaluation datasets, with significant improvements on both Qwen and LLaVA base models, demonstrating strong generalization capability. On the 3-class classification tasks (PrivateA/PrivateB), it attained Acc of 86.7% and 81.2% with corresponding F1-scores of 0.889 and 0.803, outperforming Medgamma by 15.9-24.4% in Acc. For the 2-class LIDC-IDRI benchmark, our method achieved 89.1% Acc and 0.871 F1, surpassing MedGemma by 15.9% in Acc and 0.185 in F1. The consistent superiority across different classification tasks and datasets validates the effectiveness of our nodule-focused architecture and multi-agent reasoning framework for clinical malignancy assessment.

Ablation Study. We performed an ablation study on the PrivateA dataset to assess the effectiveness of each module within our framework: Nodule Spotter, Simulated Radiologist, and DAS. It is noteworthy that the Nodule Spotter cannot be independently ablated, as the Simulated Radiologist relies on the masks it provides. As shown in **Table 3**, the results revealed a significant decline in performance across all tasks when any module was removed. This underscores the effectiveness of each component in our framework.

We examined the effectiveness of the Mask Clustering and Judge Panel components within the Nodule Spotter module, where each component builds upon the previous one. To assess their impact, we used standard object detection metrics, specifically mean Average Precision (mAP) and F1-score. As shown in **Table 4**, integrating Mask Clustering and the Judge Panel led to mAP improvements of 4% and 8%, respectively, and F1-score enhancements of 0.07 and 0.12, respectively. These results clearly demonstrate the effectiveness of both components in refining anomalous masks and improving the overall quality of mask generation.

To understand the impact of the Nodule Spotter on the overall framework’s performance, we artificially created multiple detection regions with different IoU values compared to ground truth masks. We consistently found that higher detection Acc led to better diagnostic outcomes. As illustrated in **Figure 4(a)**, these findings highlight the importance of evidence-based quantitative analysis over experience-driven qualitative assessments in multimodal diagnostic contexts, thereby affirming the effectiveness of our proposed framework.

We conducted an ablation study to assess how varying the number of medical agents influences the DAS’s Acc in grading malignancies, as illustrated in **Figure 4(b)**. The findings show that Acc increased with additional agents, peaking at five agents, after which the performance became inconsistent (Xiong et al. 2023). This trend suggests that a broader range of perspectives enhances diagnostic reasoning up to an optimal point. Beyond this threshold, the inclusion of more agents may introduce redundancy or minor inconsistencies, leading to fluctuations in performance. Consequently, we determined that deploying five medical agents strikes the ideal balance between robust diagnostic performance and compu-

Method		PrivateA				PrivateB				LIDC-IDRI			
		LLM-score	LungDLC	Pos QA	Neg QA	LLM-score	LungDLC	Pos QA	Neg QA	LLM-score	LungDLC	Pos QA	Neg QA
Generalist	GPT-4o	88.1	71.8	68.3	75.3	<u>87.2</u>	68.2	66.5	69.8	90.1	73.2	70.1	76.3
	Claude 3.7 Sonnet	86.9	65.6	58.8	72.3	86.5	64.3	55.2	73.3	88.6	65.9	59.2	72.6
	Qwen2.5-VL-7B Δ	80.2	54.3	41.4	67.2	78.2	56.3	48.2	64.5	82.5	60.3	50.1	70.4
	InternVL3-8B	77.2	56.3	43.8	68.9	75.8	54.6	41.9	67.4	77.5	61.8	54.9	68.7
	LLaVA3.2-11B \square	83.1	58.1	46.1	70.2	81.7	60.3	52.8	67.8	84.2	62.3	55.8	69.8
Medical Agent	MedGemma-27B	87.3	75.6	76.0	75.2	87.0	76.2	<u>73.7</u>	78.6	87.7	75.2	74.0	76.4
	MedR1-3B	82.1	64.5	55.4	73.5	80.5	61.1	49.9	72.3	83.3	67.2	62.7	71.7
	MedAgent-Pro	86.9	70.4	62.7	78.1	85.6	68.8	63.0	74.6	87.2	70.0	66.8	73.2
	MedAgents	83.6	65.2	55.7	74.7	82.6	63.7	54.1	73.3	84.0	68.5	61.3	75.7
	MDAgent	77.8	63.3	54.5	72.1	81.3	60.5	48.6	72.5	82.6	63.2	52.8	73.7
	LLaVA-Med	77.6	58.3	50.1	66.5	80.2	58.9	47.5	70.2	82.1	61.5	53.3	69.6
Ours	LungNoduleAgent	86.9	<u>80.5</u>	<u>77.3</u>	83.7	86.2	<u>79.5</u>	73.6	85.4	88.5	<u>83.1</u>	<u>81.6</u>	<u>84.6</u>
	Qwen2.5-VL-7B Δ	+6.7	+26.2	+35.9	+16.5	+8.0	+23.2	+25.4	+20.9	+6.0	+22.8	+31.5	+14.2
	LungNoduleAgent	<u>87.6</u>	81.9	81.6	<u>82.2</u>	89.8	80.3	76.1	<u>84.5</u>	<u>89.3</u>	83.5	82.1	84.9
	LLaVA3.2-11B \square	+4.5	+23.8	+35.5	+12.0	+8.1	+20.0	+23.3	+16.7	+5.1	21.2	+26.3	+15.1

Table 1: CT report generation results on three datasets. The best and second-best results are **bolded** and underlined, respectively. Δ , \square correspond to the base models of our LungNoduleAgent for comparisons, and + shows the improvements between them.

Method		PrivateA		PrivateB		LIDC-IDRI	
		Acc	F1	Acc	F1	Acc	F1
Gen.	GPT-4o	46.2	39.7	41.2	56.2	64.1	76.7
	Claude 3.7 Sonnet	33.1	35.6	34.2	39.8	62.3	68.7
	Qwen2.5-VL-7B Δ	31.4	35.4	31.1	36.8	56.5	60.2
	InternVL3-8B	36.1	39.8	35.3	44.6	58.3	60.8
	LLaVA3.2-11B \square	35.6	40.3	36.7	45.2	59.1	61.2
	Med.	MedGemma-27B	62.3	68.3	60.2	70.6	73.2
MedR1-3B		52.3	62.7	48.9	58.6	67.6	73.5
MedAgent-Pro		60.2	65.9	60.1	71.9	72.6	74.1
MedAgents		54.3	64.2	52.3	60.1	65.3	58.2
MDAgent		45.6	52.1	43.9	56.5	66.7	63.2
LLaVA-Med		39.3	38.2	40.7	45.9	62.5	65.8
Ours	LungNoduleAgent	<u>85.8</u>	<u>82.1</u>	<u>80.6</u>	<u>75.6</u>	<u>88.5</u>	<u>84.1</u>
	Qwen2.5-VL-7B Δ	+54.4	+46.7	+49.5	+38.8	+32.0	+23.9
	LungNoduleAgent	86.7	88.9	81.2	80.3	89.1	87.1
	LLaVA3.2-11B \square	+51.1	+48.6	+44.5	+35.1	+30.0	+25.9

Table 2: Malignancy prediction results on three datasets. The best and second-best results are **bolded** and underlined, respectively. Δ , \square correspond to the base models of our LungNoduleAgent for comparisons, and + shows the improvements between them.

NS	SR	DAS	Acc(%)	LungDLC
-	-	✓	62.1	57.9
✓	✓	-	66.7	88.9
✓	-	✓	<u>75.1</u>	<u>67.3</u>
✓	✓	✓	86.7	88.9

Table 3: Ablation study on different module combinations. The best and second-best results are **bolded** and underlined, respectively. Abbreviations: Nodule Spotter (NS), Simulated Radiologist (SR), Doctor Agent System (DAS).

tational efficiency in our system’s analyses.

MoE	Clustering	Judge Panel	mAP(%)	F1(10 ⁻²)
✓	-	-	67.1	64.3
✓	✓	-	<u>71.6</u>	<u>71.3</u>
✓	✓	✓	79.3	83.5

Table 4: Ablation study on the effect of MoE, Clustering, and Judge Panel. The best and second-best results are **bolded** and underlined, respectively.

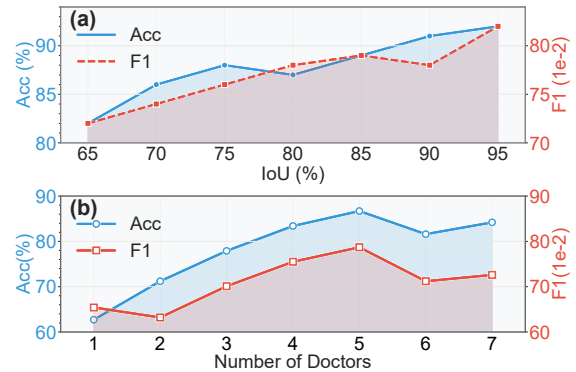


Figure 4: (a): Ablation study on the influence of nodule detection Acc on malignancy grading. (b): Ablation study on the impact of the number of medical agents in the DAS on nodule malignancy grading.

Conclusion

We introduce LungNoduleAgent, a novel multi-agent system designed to mimic the clinical workflow of lung nodule diagnosis in CT scans. Our framework is specifically tailored to produce accurate malignancy grading and detailed CT reports, proving the significance of region-level semantic alignment and multi-agent collaboration in diagnosing nodules. Compared to existing VLMs and medical agents, our framework shows superior performance and remarkable generalization capabilities. Ablation studies further validate the effectiveness of each component within the system.

Acknowledgements

This work was supported by National Natural Science Foundation of China (No.62076084, 61702146, U20A20386, U22A2033) and Zhejiang Provincial Natural Science Foundation of China (No. LY21F020017, 2023C03090), Guangxi Science and Technology Program (No. FN2504240022), the Guangxi Key R&D Project (No. AB24010167), the Project (No. 20232ABC03A25), GuangDong Basic and Applied Basic Research Foundation (No. 2025A1515011617, 2022A1515110570), Futian Healthcare Research Project (No.FTWS002), Central Funds Guiding the Local Science and Technology Development Project (No. 2025ZYDF106).

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2024. Claude 3.7 Sonnet.
- Arslan, M.; Ghanem, H.; Munawar, S.; and Cruz, C. 2024. A Survey on RAG with LLMs. *Procedia Computer Science*, 246: 3781–3790.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.
- Cao, H.; Liu, H.; Song, E.; Ma, G.; Xu, X.; Jin, R.; Liu, T.; and Hung, C.-C. 2020. A two-stage convolutional neural networks for lung nodule detection. *IEEE Journal of Biomedical and Health Informatics*, 24(7): 2006–2015.
- Chan, C.-M.; Chen, W.; Su, Y.; Yu, J.; Xue, W.; Zhang, S.; Fu, J.; and Liu, Z. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.
- Driessen, R.; Nandwana, S.; Hajibonabi, F.; Moreno, C.; Davarpanah, A.; and Balthazar, P. 2025. Completeness and Accuracy of Malignancy History in Abdominal CT Order Requisitions and Final Radiology Reports. *Current Problems in Diagnostic Radiology*.
- Du, Y.; Li, S.; Torralba, A.; Tenenbaum, J. B.; and Mordatch, I. 2023. Improving factuality and reasoning in language models through multiagent debate. In *Forty-first International Conference on Machine Learning*, 11733–11763.
- Edge, D.; Trinh, H.; Cheng, N.; Bradley, J.; Chao, A.; Mody, A.; Truitt, S.; Metropolitansky, D.; Ness, R. O.; and Larson, J. 2024. From local to global: A graph rag approach to query-focused summarization. *arXiv preprint arXiv:2404.16130*.
- Fallahpour, A.; Ma, J.; Munim, A.; Lyu, H.; and Wang, B. 2025. Medrax: Medical reasoning agent for chest x-ray. *arXiv preprint arXiv:2502.02673*.
- Fan, C.; Liu, L.; Wang, Y.; Li, D.; Liang, Q.; Elazab, A.; Liu, Z.; Hu, J.; Tian, Y.; Zhang, Y.; et al. 2025. Deep Neural Network for Lung Adenocarcinoma Subtype from Multimodal Fusion of Imaging and Clinical Data. In *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Hammer, M. M.; Palazzo, L. L.; Eckel, A. L.; Barbosa Jr, E. M.; and Kong, C. Y. 2019. A decision analysis of follow-up and treatment algorithms for nonsolid pulmonary nodules. *Radiology*, 290(2): 506–513.
- Ji, Z.; Wu, Y.; Zeng, X.; An, Y.; Zhao, L.; Wang, Z.; and Ganchev, I. 2023. Lung nodule detection in medical images based on improved YOLOv5s. *IEEE Access*, 11: 76371–76387.
- Kim, Y.; Park, C.; Jeong, H.; Chan, Y. S.; Xu, X.; McDuff, D.; Lee, H.; Ghassemi, M.; Breazeal, C.; and Park, H. W. 2024. Mdagents: An adaptive collaboration of llms for medical decision-making. *Advances in Neural Information Processing Systems*, 37: 79410–79452.
- Ladbury, C.; Amini, A.; Govindarajan, A.; et al. 2023. Integration of artificial intelligence in lung cancer: Rise of the machine. *Cell Reports Medicine*, 4(2).
- Lai, Y.; Zhong, J.; Li, M.; Zhao, S.; and Yang, X. 2025. Med-r1: Reinforcement learning for generalizable medical reasoning in vision-language models. *arXiv preprint arXiv:2503.13939*.
- Lee, C.; Park, S.; Shin, C.-I.; Choi, W. H.; Park, H. J.; Lee, J. E.; and Ye, J. C. 2024. Read like a radiologist: Efficient vision-language model for 3d medical imaging interpretation. *arXiv preprint arXiv:2412.13558*.
- Li, B.; Yan, T.; Pan, Y.; Luo, J.; Ji, R.; Ding, J.; Xu, Z.; Liu, S.; Dong, H.; Lin, Z.; et al. 2024a. Mmedagent: Learning to use medical tools with multi-modal agent. *arXiv preprint arXiv:2407.02483*.
- Li, C.; Wong, C.; Zhang, S.; Usuyama, N.; Liu, H.; Yang, J.; Naumann, T.; Poon, H.; and Gao, J. 2023. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *arXiv preprint arXiv:2306.00890*.
- Li, D.; Jiang, B.; Huang, L.; Beigi, A.; Zhao, C.; Tan, Z.; Bhattacharjee, A.; Jiang, Y.; Chen, C.; Wu, T.; Shu, K.; Cheng, L.; and Liu, H. 2024b. From Generation to Judgment: Opportunities and Challenges of LLM-as-a-judge. *arXiv preprint arXiv: 2411.16594*.
- Li, D.; Sun, R.; Huang, Y.; Zhong, M.; Jiang, B.; Han, J.; Zhang, X.; Wang, W.; and Liu, H. 2025. Preference leakage: A contamination problem in llm-as-a-judge. *arXiv preprint arXiv:2502.01534*.
- Liang, T.; He, Z.; Jiao, W.; Wang, X.; Wang, Y.; Wang, R.; Yang, Y.; Shi, S.; and Tu, Z. 2023. Encouraging divergent thinking in large language models through multi-agent debate. *arXiv preprint arXiv:2305.19118*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023a. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36: 34892–34916.
- Liu, J.; Wang, W.; Ma, Z.; Huang, G.; SU, Y.; Chang, K.-J.; Chen, W.; Li, H.; Shen, L.; and Lyu, M. 2024. Medchain: Bridging the gap between llm agents and clinical practice through interactive sequential benchmarking. *arXiv preprint arXiv:2412.01605*.

- Liu, X.; Yu, H.; Zhang, H.; Xu, Y.; Lei, X.; Lai, H.; Gu, Y.; Ding, H.; Men, K.; Yang, K.; et al. 2023b. Agentbench: Evaluating llms as agents. *arXiv preprint arXiv:2308.03688*.
- Mao, Y.; Xu, W.; Qin, Y.; and Gao, Y. 2025. CT-Agent: A Multimodal-LLM Agent for 3D CT Radiology Question Answering. *arXiv preprint arXiv:2505.16229*.
- Martinez-Murcia, F. J.; Ortiz, A.; Ramírez, J.; Górriz, J. M.; and Cruz, R. 2021. Deep residual transfer learning for automatic diagnosis and grading of diabetic retinopathy. *Neurocomputing*, 452: 424–434.
- Nakach, F.-Z.; Idri, A.; and Goceri, E. 2024. A comprehensive investigation of multimodal deep learning fusion strategies for breast cancer classification. *Artificial Intelligence Review*, 57(12): 327.
- Osarogiagbon, R. U.; Liao, W.; Faris, N. R.; Fehnel, C.; Goss, J.; Shepherd, C. J.; Qureshi, T.; Matthews, A. T.; Smeltzer, M. P.; and Pinsky, P. F. 2023. Evaluation of lung cancer risk among persons undergoing screening or guideline-concordant monitoring of lung nodules in the Mississippi Delta. *JAMA Network Open*, 6(2): e230787–e230787.
- Qiu, Y.; Chen, X.; Wu, X.; Li, Y.; Xu, P.; Jin, K.; Shang, X.; Chotcomwongse, P.; He, M.; and Shi, D. 2025. Embodied artificial intelligence in ophthalmology. *NPJ Digital Medicine*, 8(1): 351.
- Raza, R.; Zulfiqar, F.; Khan, M. O.; Arif, M.; Alvi, A.; Iftikhar, M. A.; and Alam, T. 2023. Lung-EffNet: Lung cancer classification using EfficientNet from CT-scan images. *Engineering Applications of Artificial Intelligence*, 126: 106902.
- Samuel, G. 2011. The Lung Image Database Consortium (LIDC) and Image Database resource initiative (IDRI): A completed reference database of lung nodules on CT scans. *Medical Physics*, 38: 2.
- Sellergren, A.; Kazemzadeh, S.; Jaroensri, T.; Kiraly, A.; Traverse, M.; Kohlberger, T.; Xu, S.; Jamil, F.; Hughes, C.; Lau, C.; et al. 2025. MedGemma Technical Report. *arXiv preprint arXiv:2507.05201*.
- Shen, Y.; Fang, Z.; Zhuang, K.; Zhou, G.; Yu, X.; Zhao, Y.; Tian, Y.; Ge, R.; Wang, C.; Fan, X.; et al. 2025. CSF-NET: Cross-Modal Spatiotemporal Fusion Network for Pulmonary Nodule Malignancy Predicting. In *2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Swanson, K.; Wu, E.; Zhang, A.; Alizadeh, A. A.; and Zou, J. 2023. From patterns to patients: Advances in clinical machine learning for cancer diagnosis, prognosis, and treatment. *Cell*, 186(8): 1772–1791.
- Tammemagi, M.; Ritchie, A. J.; Atkar-Khattra, S.; Dougherty, B.; Sanghera, C.; Mayo, J. R.; Yuan, R.; Manos, D.; McWilliams, A. M.; Schmidt, H.; et al. 2019. Predicting malignancy risk of screen-detected lung nodules—mean diameter or volume. *Journal of Thoracic Oncology*, 14(2): 203–211.
- Tang, X.; Zou, A.; Zhang, Z.; Li, Z.; Zhao, Y.; Zhang, X.; Cohan, A.; and Gerstein, M. 2023. Medagents: Large language models as collaborators for zero-shot medical reasoning. *arXiv preprint arXiv:2311.10537*.
- Thirunavukarasu, A. J.; Ting, D. S. J.; Elangovan, K.; Gutierrez, L.; Tan, T. F.; and Ting, D. S. W. 2023. Large language models in medicine. *Nature Medicine*, 29(8): 1930–1940.
- UrRehman, Z.; Qiang, Y.; Wang, L.; Shi, Y.; Yang, Q.; Khatkhat, S. U.; Aftab, R.; and Zhao, J. 2024. Effective lung nodule detection using deep CNN with dual attention mechanisms. *Scientific Reports*, 14(1): 3934.
- Wang, C.; Elazab, A.; Jia, F.; Wu, J.; and Hu, Q. 2018. Automated chest screening based on a hybrid model of transfer learning and convolutional sparse denoising autoencoder. *Biomedical Engineering Online*, 17(1): 63.
- Wang, H.; Zhao, S.; Qiang, Z.; Xi, N.; Qin, B.; and Liu, T. 2024. Beyond direct diagnosis: LLM-based multi-specialist agent consultation for automatic diagnosis. *arXiv preprint arXiv:2401.16107*.
- Wang, Z.; Wu, J.; Cai, L.; Low, C. H.; Yang, X.; Li, Q.; and Jin, Y. 2025. MedAgent-Pro: Towards Evidence-Based Multi-Modal Medical Diagnosis via Reasoning Agentic Workflow. *arXiv preprint arXiv:2503.18968*.
- Xiao, X.; Zhang, Y.; Nguyen, T.-H.; Lam, B.-T.; Wang, J.; Zhao, L.; Hamm, J.; Wang, T.; Li, X.; Wang, X.; et al. 2025. Describe Anything in Medical Images. *arXiv preprint arXiv:2505.05804*.
- Xiong, K.; Ding, X.; Cao, Y.; Liu, T.; and Qin, B. 2023. Examining inter-consistency of large language models collaboration: An in-depth analysis via debate. *arXiv preprint arXiv:2305.11595*.
- Yang, C.; Lu, J.; Wan, H.; Yu, J.; and Qin, F. 2025. From What to Why: A Multi-Agent System for Evidence-based Chemical Reaction Condition Reasoning. *arXiv preprint arXiv:2509.23768*.
- Yu, X.; Chen, Z.; Zhang, Y.; Lu, S.; Shen, R.; Zhang, J.; Hu, X.; Fu, Y.; and Yan, S. 2025. Visual document understanding and question answering: A multi-agent collaboration framework with test-time scaling. *arXiv preprint arXiv:2508.03404*.
- Zhang, D.; Yu, Y.; Dong, J.; Li, C.; Su, D.; Chu, C.; and Yu, D. 2024. Mm-llms: Recent advances in multimodal large language models. *arXiv preprint arXiv:2401.13601*.
- Zhang, K.; Li, D.; Luo, W.; and Ren, W. 2021. Dual attention-in-attention model for joint rain streak and rain-drop removal. *IEEE Transactions on Image Processing*, 30: 7608–7619.
- Zhang, X.; Wu, C.; Zhao, Z.; Lin, W.; Zhang, Y.; Wang, Y.; and Xie, W. 2023. Pmc-vqa: Visual instruction tuning for medical visual question answering. *arXiv preprint arXiv:2305.10415*.
- Zhu, J.; Wang, W.; Chen, Z.; Liu, Z.; Ye, S.; Gu, L.; Tian, H.; Duan, Y.; Su, W.; Shao, J.; et al. 2025. Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.