

MedSyn: Enhancing Diagnostics with Human-AI Collaboration

Burcu Sayin^{1,*}, Ipek Baris Schlicht², Ngoc Vo Hong³, Sara Allievi³, Jacopo Staiano¹, Pasquale Minervini⁴ and Andrea Passerini¹

¹University of Trento, Italy

²Universitat Politècnica de València, Spain

³Santa Chiara Hospital, Trento, Italy

⁴The University of Edinburgh, UK

Abstract

Clinical decision-making is inherently complex, often influenced by cognitive biases, incomplete information, and case ambiguity. Large Language Models (LLMs) have shown promise as tools for supporting clinical decision-making, yet their typical one-shot or limited-interaction usage may overlook the complexities of real-world medical practice. In this work, we propose a hybrid human-AI framework, MedSyn, where physicians and LLMs engage in multi-step, interactive dialogues to refine diagnoses and treatment decisions. Unlike static decision-support tools, MedSyn enables dynamic exchanges, allowing physicians to challenge LLM suggestions while the LLM highlights alternative perspectives. Through simulated physician-LLM interactions, we assess the potential of open-source LLMs as physician assistants. Results show open-source LLMs are promising as physician assistants in the real world. Future work will involve real physician interactions to further validate MedSyn's usefulness in diagnostic accuracy and patient outcomes.

Keywords

hybrid decision making, medical decision making, hybrid intelligence, clinical NLP, LLM agents

1. Introduction

In traditional clinical practice, a physician's diagnosis and treatment plan may be influenced by cognitive biases, incomplete information, or the inherent complexity of the case [1, 2]. Additionally, physicians often work in time-sensitive, high-pressure environments (e.g., emergency departments), where cognitive overload can increase the risk of misdiagnosis. Recent advancements in Large Language Models (LLMs) offer new opportunities for AI-assisted medical decision-making [3, 4, 5, 6]. We propose that physicians and LLMs can effectively cooperate within multi-step interactive scenarios wherein the LLM's suggestions – whether accurate or flawed – serve as opportunities for deeper inquiry and reflection. Thus, in this work, we investigate to what extent such a hybrid cooperative human-AI setup allows physicians to uncover potential oversights, recognize overlooked symptoms, and reconsider treatment options. Unlike static systems that provide one-time recommendations, we propose a dynamic conversational framework that evolves based on real-time interactions, ensuring that physicians maintain control over the clinical decision-making process. Specifically, we explore the collaboration of physicians and LLMs on a specific and sensitive topic: a patient's diagnosis. For instance, if the physician overlooks key

HHAI-WS 2025: Workshops at the Fourth International Conference on Hybrid Human-Artificial Intelligence (HHAI), June 9–13, 2025, Pisa, Italy

*Corresponding author.

✉ burcu.sayin@unitn.it (B. Sayin); ibarsch@doctor.upv.es (I. B. Schlicht); ngoc.vohong@apss.tn.it (N. V. Hong); sara.allievi@apss.tn.it (S. Allievi); jacopo.staiano@unitn.it (J. Staiano); p.minervini@ed.ac.uk (P. Minervini); andrea.passerini@unitn.it (A. Passerini)

🌐 <https://burcusayin.github.io/> (B. Sayin); <https://www.apss.tn.it/Azienda/Dirigenti/Vo-Hong-Ngoc> (N. V. Hong);

<https://www.apss.tn.it/Azienda/Dirigenti/Allievi-Sara> (S. Allievi); <https://www.staiano.net/> (J. Staiano);

<http://www.neuralnoise.com/> (P. Minervini); <https://disi.unitn.it/~passerini/index.html> (A. Passerini)

🆔 0000-0001-6804-127X (B. Sayin); 0000-0002-5037-2203 (I. B. Schlicht); 0009-0001-4641-0190 (N. V. Hong);

0000-0003-4428-9494 (S. Allievi); 0000-0002-1260-4640 (J. Staiano); 0000-0002-8442-602X (P. Minervini); 0000-0002-2765-5395 (A. Passerini)



© 2025 Copyright for this paper by its authors.

Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).




Figure 1: MedSyn Framework

symptoms or suggests a suboptimal treatment, the LLM can ask patient-specific follow-up questions or recommend reconsidering the diagnosis. Conversely, if an LLM proposes an incorrect diagnosis, the physician can critically examine its reasoning, prompting the model to refine its suggestion. This iterative exchange improves diagnostic accuracy and therapeutic decision-making, serving as a cognitive safety net that aids physicians in complex, ambiguous cases with a higher risk of error.

This working paper presents our initial efforts on building **MedSyn**, a medical synergy framework that positions LLMs as conversational partners in clinical decision-making. By fostering human-AI collaboration, MedSyn aims to enhance diagnostics while preserving the physician’s critical role in patient care. To evaluate MedSyn, we curate and merge data from MIMIC-IV [7] and MIMIC-IV-Note [8, 9], creating a diverse set of patient records for model assessment. We then investigate 25 open-source chat-based and medical-domain LLMs to evaluate their capacity for multi-turn engagement. Our analysis highlights both the challenges and opportunities in developing open-source medical dialogue systems. While several models struggled to maintain coherent, multi-turn interactions, others demonstrated the ability to engage in sustained, in-depth discussions about patient conditions. From the 25 evaluated models, we selected three promising candidates for further experimentation—LLaMA3 (8B and 70B) [10] and Gemma2 (27B) [11]. We also included DeepSeek-R1 [12], distilled to Llama3.3-70B-Instruct available via Ollama¹, as a representative of state-of-the-art open-source models that currently fall short in handling complex medical multi-turn dialogues. To assess the role of iterative questioning and collaborative reasoning, we simulate physician–LLM conversations in a controlled setting. Preliminary results show that interactive, multi-step exchanges yield more comprehensive patient assessments and enhance diagnostic clarity. These findings are qualitatively supported by physician analysis of both LLM decisions and their corresponding dialogue traces. As a next step, we aim to replace the simulated physician LLM with real clinicians, enabling direct interaction with the assistant LLM. This will help refine MedSyn for clinical deployment and further validate its utility in real-world medical settings.

2. MedSyn

Figure 1 shows the overview of the MedSyn framework. MedSyn receives the clinical note of the patient as input. The clinical note includes several information, including the chief complaint, history of the present illness, physical exam, and pertinent results. Given the limited time physicians have for each patient in the real world, it can be challenging for them to thoroughly review all the details in a clinical note, analyze the patient’s condition, and integrate their observations to make an accurate diagnosis. The MedSyn framework supports physicians through an LLM-based virtual assistant that has access to

¹<https://ollama.com/library/deepseek-r1:70b>

all the details in the clinical note, while the physician is assumed to have access only to the patient’s chief complaint. To gather necessary information about the patient and engage in a collaborative discussion, the physician initiates a multi-turn interaction. In the first turn, the physician asks the assistant for an initial evaluation of the patient. In response, the assistant carefully analyzes the clinical note and provides a detailed observation. Following this, the physician and the virtual assistant engage in a dynamic discussion about the patient’s condition. This exchange continues until the physician feels they have gathered all the necessary information and is confident in their understanding of the patient’s condition. At this point, the physician concludes the discussion and drafts the discharge text for the patient. The discharge text may include several sections, such as the discharge diagnosis, condition, medications, and the follow-up instructions. For this study, however, we focus solely on the “diagnosis” and the corresponding “ICD-10 codes”² used by clinicians to code and classify medical diagnoses.

3. Experimental Work

We combined MIMIC-IV³ [7, 9] and MIMIC-IV-Note⁴ [8] datasets by selecting records with ICD-10 coding [13], which covers diseases from coarse, “chapter” level (e.g. E00-E90) to finer granularities (e.g. E10.9 where E10 is a disease category and 10.9 indicates the disease code). The resulting merged dataset contained 122,266 records spanning 5,802 unique diagnoses. Upon analyzing the discharge text field in these records, we observed that most followed a common structure, though certain subsections varied (e.g., the “major surgical or invasive procedure” section was present in some records but absent in others). Samples with missing headings or free-form discharge notes hindered effective parsing and prevented the establishment of a standardized format across all records. After consulting with three physicians, we identified the most important sections for our experiments and excluded samples that did not conform to the expected format. Specifically, we selected records that include the following sections in their discharge texts: “chief complaint, history of present illness, social history, physical exam, pertinent results, major surgical or invasive procedure, brief hospital course, medications on admission, discharge medications, discharge diagnosis, discharge condition, and discharge instructions”. Furthermore, we removed records where the patient’s status was “deceased” or “expired”. This filtering process resulted in a final dataset of 74,850 records. Then, we randomly (seed=13) selected 1,000 records as our test set. It consists of 2,350 unique diagnoses (on a total of 13,384). The average number of ICD-10 codes appearing in a sample is 5.61. The most common diagnosis is ‘E78.5’ (Hyperlipidemia, Unspecified), while 1,112 diagnoses are identified as the rarest (e.g. ‘H53.40’: Unspecified visual field defects). Since access to this dataset requires completing specialized training, *CITI*,⁵ we are unable to publicly share our test set and LLM outputs. However, we have detailed our preprocessing steps above and made our code available.⁶

3.1. Models & Frameworks

We investigated 25 open-source models⁷ across general-purpose, chat-based, and medical domains, finding that most struggled with multi-turn dialogues. Some chat-based models (e.g., OpenChat:7B [14]) performed poorly in medical conversations, while certain medical domain models (e.g., Meditron:7B [15] and MedLlama2:7B)⁸ exhibited limitations in handling real-world dialogues. Among the evaluated models, we identified three promising candidates within our experimental setup: Llama3 (8B and 70B) [10] and Gemma2:27B [11]. To illustrate the challenges even state-of-the-art models face in medical

²<https://www.icd10data.com/ICD10CM/Codes/>

³<https://physionet.org/content/mimiciv/3.0/>

⁴<https://physionet.org/content/mimic-iv-note/2.2/>

⁵<https://physionet.org/about/citi-course/>

⁶See our source code here: <https://github.com/burcusayin/MedSyn>

⁷command-r-plus:104b, command-r:35b, openchat:7b, mistral:7b, mistral:lite:7b, mixtral:8x7b, qwen2:7b, meditron:7b, meditron:70b, medllama2:7b, llama3-chatqa:8b and 70b, llama3:8b and 70b, llama3.1:8b, llama3.2:3b, dolphin-llama3:8b, dolphin-llama3:70b, phi3:14b, nemotron:70b, alfred:40b, deepseek-R1-Distill-Llama-70B, tulu3:8b and 70b, gemma2:27b

⁸https://huggingface.co/llSourceCell/medllama2_7b

dialogues, we present results with DeepSeek-R1:70B [12] (Distilled to Llama-70B, available by Ollama⁹). We implemented our multi-agent environment using Ollama¹⁰ and Langroid.¹¹

3.2. Use cases

To assess the potential of our framework for real-world deployment in medical decision-making systems, we simulated interactions using LLMs—one serving as the chief physician and another as the physician assistant. As a baseline, we defined the “phy w/complaint” scenario, in which the physician LLM receives only the patient’s chief complaint from the clinical note and generates the discharge text without any interaction or dialogue. In contrast, the “two agent” setup simulates the collaboration between physicians and assistants in the real world by implementing the MedSyn pipeline (Section 2). Here, the physician agent is limited to the chief complaint, while the assistant agent has access to the complete clinical note, including the history of present illness, physical examination, and pertinent results. Both configurations employ zero-shot prompting, with full prompt details provided below.

Baseline Case We use the baseline prompt in the “phy w/complaint” case.

Baseline Prompt

You are Dr. Ellis, the chief physician responsible for reviewing clinical notes and writing a discharge text for patients.

****Here is the clinical note for the patient:**** {clinicalNote}.

Instructions:

1. Carefully analyze the given clinical note to ensure that no symptoms are overlooked.
2. You are not allowed to ask any questions or make assumptions beyond the information provided in the clinical note.
3. Once you are ready, write the discharge text for the patient.
4. The discharge text should include only the ‘diagnosis’ and ‘codes’ fields:
 - ‘diagnosis’ field should specify the patient’s final diagnosis. Please note that you should decide the final diagnosis.
 - ‘codes’ field should list the ICD-10 codes corresponding to the diagnosis specified in the ‘diagnosis’ field. Keep in mind that this field is a string, do not use ‘[]’ while listing the codes.
5. Remember to refer the clinical note while writing the discharge text. Ensure that the ‘diagnosis’, and ‘codes’ fields are complete and unambiguous; they must not be left empty or unclear.
6. Return your dischargeText using the TOOL ‘baseline_discharge_text_tool’.

Two-agent Case We use different prompts for the chief physician and physician assistant LLMs.

Chief Physician Prompt

You are Dr. Ellis, the Chief Physician, collaborating with Dr. Lee, your assistant. Your task is to review a clinical note by initiating an evaluation from Dr. Lee and engaging in a natural, focused conversation to assess the patient’s condition. Avoid fabricating interactions or simulating dialogue with Dr. Lee. Instead, clearly articulate your questions or follow-ups, analyze Dr. Lee’s responses, and use this information to guide your decision-making.

Your responsibilities include the following:

- Verify the patient’s condition, symptoms, and diagnosis.
- Ensure all symptoms are accounted for and thoroughly understand the patient’s condition to deliver optimal care.
- Address doubts regarding the diagnosis or treatment plan by conducting further evaluations with Dr. Lee to achieve accurate and effective results.

⁹<https://ollama.com/library/deepseek-r1>

¹⁰<https://github.com/ollama>

¹¹<https://github.com/langroid/langroid>

****Here is the clinical note for the patient:**** {clinicalNote}.

Instructions:

1. Begin by requesting an initial evaluation of the patient from Dr. Lee.
2. Engage in a collaborative discussion to confirm the patient's diagnosis. Please note that Dr. Lee has access to a more detailed clinical note, so you **MUST** consult to Dr. Lee to obtain the necessary information for making the diagnosis.
3. Keep in mind that you have limited time for every patient. Please avoid duplicate recommendations, conversations, and questions related to treatments. Keep each message **CONCISE** and to the point.
4. Once you have gathered sufficient information and are confident in the diagnosis, stop the discussion and write the patient's discharge text.
5. The discharge text should include only the 'diagnosis' and 'codes' fields:
 - 'diagnosis' field should specify the patient's final diagnosis. Please note that you should decide the final diagnosis.
 - 'codes' field should list the ICD-10 codes corresponding to the diagnosis specified in the 'diagnosis' field.
6. Remember to refer to your discussion with Dr. Lee and the clinical note while writing the discharge text. Ensure that the 'diagnosis', and 'codes' fields are complete and unambiguous; they must not be left empty or unclear.
7. Do **NOT** ask Dr. Lee to check or write your dischargeText. It is **YOUR RESPONSIBILITY** to write and submit the dischargeText.
8. Return your dischargeText using the TOOL 'discharge_text_tool'. Do **NOT** mention the TOOL 'discharge_text_tool' to Dr. Lee.

Physician Assistant Prompt

You are Dr. Lee, an assistant physician working under the supervision of Dr. Ellis, the chief physician. Your role is to review the patient's clinical notes to perform an initial evaluation, which will support Dr. Ellis in assessing the patient's condition and writing the discharge text. Following your evaluation, you will engage in a collaborative discussion with Dr. Ellis to confirm the diagnosis and determine the next steps.

****Here is the clinical note for the patient:**** {clinicalNote}.

Task: Thoroughly analyze the clinical note and provide a structured summary that includes:

- Key symptoms: Highlight notable symptoms that may require further investigation.
- Preliminary diagnosis: Offer an initial diagnosis based on the patient's symptoms and medical history.
- Potential complications: Identify any critical issues or risks Dr. Ellis should consider.
- Recommendations: Suggest further evaluations if uncertainties remain about the patient's condition.

Instructions:

1. Ensure your evaluation is clear, precise, and structured to facilitate an informed discussion.
2. In each round of the discussion, limit yourself to a **CONCISE** message.
3. Keep in mind that you have limited time for every patient. Please avoid duplicate recommendations, conversations, and questions related to treatments.

Process: You will first receive a message from Dr. Ellis, asking for your initial assessment. Afterward, you can follow up in each discussion round to collaboratively refine the diagnosis.

4. Results

Directly comparing discharge texts with LLM responses using standard metrics presents several challenges: (i) Discharge texts lack the conversational tone of LLM responses, (ii) LLMs may generate

variable lengths of ICD-10 codes and diagnoses, including occasional hallucinated codes,¹² (iii) Physicians often employ abbreviations and specialized formatting in discharge texts, whereas LLMs produce more standard, conversational sentences, and (iv) The ground truth for diagnoses and ICD-10 codes is longer than LLM outputs. According to two physicians from the in-house annotators, this discrepancy arises because physicians include codes for current and past illnesses based on system recommendations, while LLMs are limited to the information provided in the prompt, which in our case focuses on current symptoms rather than a comprehensive patient history. Thus, specific metrics designed for ICD code detection [16] are unsuitable.

ICD-10 Classification As stated in §3, ICD-10 contains coarse and fine-grained definitions of diseases. In preliminary experiments, we observed that all LLMs tended to not generate fine-grained codes, which could be expected in our zero-shot multi-label classification setup. We explored this issue by discussing several ground truth examples with physicians: they brought to our attention that when selecting ICD-10 subcodes – often very specific to the diagnosis – different physicians might choose different codes among those corresponding to the same primary diagnosis; most importantly, it was highlighted how physicians tend to include codes for all the acute or chronic conditions a patient is affected in the patient’s medical record, hence including several codes actually unrelated to the specific chief complaint. This characteristic of the ground truth makes the selection of evaluation metrics challenging, as it is impossible to selectively remove the ICD codes unrelated to the chief complaint. For this reason, we resort to compute precision, recall, F1-score, and Jaccard similarity score¹³ on a per-sample basis, and report the mean values in Table 1. F1 and Recall show that the agents struggled to accurately predict disease categories, frequently missing ICD codes present in the ground truth. Regarding Precision, all models performed better in predicting disease chapters, a simpler task than detecting disease categories. DeepSeek-R1 and Llama3:70B performed best in the “phy w/complaint” case (in terms of precision), with the former excelling in Disease Category and the latter in Disease Chapter.

In two-agent case, we observed that DeepSeek-R1 struggled to engage in dialogue. Despite explicitly stating in the prompt that it *must* consult the assistant before making a diagnosis, it often relied on internal reasoning and directly generated the discharge text, with minimal interaction with its assistant. Figure 2 shows the number of turns each <chief physician agent, Llama3:8B> pair produced per sample in the “two-agent” case. Notably, DeepSeek-R1:70B engaged in conversations infrequently, whereas Llama:70B exhibited higher interaction, averaging 19.2 turns per sample. Both the Llama3:70B and Gemma2:27B models demonstrated strong performance in engaging in effective dialogues with their assistants and generating well-structured discharge summaries. However, Gemma2:27B was more effective in dialogues, generating the discharge text in 9.3 turns in average. Additionally, Llama3:8B proved to be an effective physician assistant by responding concisely to the chief physician and extracting the necessary information from the clinical note. This is evident from their performance, which closely approaches the performance in “phy w/full_note” case and generates the discharge text without any interaction with the assistant. Our preliminary findings suggest that open-source LLMs hold promise as physician assistants in real-world clinical settings. However, further analysis needed to clarify the limitations and improve performance.

Qualitative Analysis by Physicians The use of LLMs in a healthcare setting has shown interesting results from a clinical perspective. The “phy w/complaint” case showed that, starting from the main symptom, LLM was able to identify a possible diagnosis despite having no access to additional clinical and instrumental information. However, it could only align with a subset of the physician’s diagnostic hypothesis and was unable to provide a detailed diagnosis. On the other hand, the “two-agent” scenario yielded better results in terms of diagnostic precision and completeness. In particular, the Gemma2:27B model made precise diagnoses when interacted with the Llama3:8B model, identifying even rare

¹²For instance, writing the code M3459 for diagnosis “Multiple Sclerosis Flare”: the code M3459 does not exist; “M34” corresponds to “systemic sclerosis” disease which is unrelated to “multiple sclerosis” (“G35”).

¹³Please see our code for the evaluation: <https://github.com/burcusaayin/MedSyn/blob/main/src/evaluation/metrics.py>

Table 1

Performance of Llama3:70B, Gemma2:27B, and DeepSeekR1:70B models as chief physician in ICD-10 disease category and chapter prediction. We use Llama3:8B as the physician assistant. “phy w/full_note” refers to the reference performance of chief physician agents when given full access to the clinical note, as opposed to only the chief complaint, without any interaction with Llama3:8B. We compare the performance in the “phy w/complaint” and “two-agent” cases, highlighting the best-performing ones.

Agent	Case	Disease Category				Disease Chapter			
		Jaccard	Precision	Recall	F1	Jaccard	Precision	Recall	F1
Llama3:70B	phy w/complaint	0.04	0.3	0.04	0.07	0.15	0.75	0.15	0.23
	two-agent	0.07	0.37	0.08	0.12	0.22	0.82	0.23	0.34
	phy w/full_note	0.09	0.35	0.1	0.14	0.23	0.78	0.25	0.35
Gemma2:27B	phy w/complaint	0.03	0.26	0.04	0.1	0.14	0.7	0.14	0.22
	two-agent	0.06	0.38	0.06	0.1	0.18	0.82	0.18	0.28
	phy w/full_note	0.06	0.34	0.1	0.11	0.19	0.8	0.2	0.3
DeepSeekR1:70B	phy w/complaint	0.04	0.31	0.04	0.07	0.13	0.72	0.13	0.21
	two-agent	0.04	0.29	0.04	0.06	0.14	0.74	0.14	0.23
	phy w/full_note	0.09	0.5	0.09	0.14	0.2	0.81	0.21	0.31




Figure 2: Histogram of multi-turn interactions across physician agents, each engaging with Llama3:8B. DeepSeek-R1 rarely engaged in dialogues. Llama3:70B and Gemma2:27B demonstrated effective interactions.

conditions that could be overlooked by a physician (e.g., Ludwig’s angina). The interaction between the physician LLM and the assistant LLM allowed for a more complete diagnosis, as the physician could obtain additional information regarding the patient’s characteristics and instrumental exams. In this case, the main challenge was distinguishing between acute and chronic conditions, as there were instances where the chief physician agent identified a pre-existing condition as the primary diagnosis. DeepSeek-R1 did not perform well in “two-agent” case, and did not improve the diagnosis compared to “phy w/complaint” case, often merely repeating the diagnosis already made. Regarding the identification of ICD-10 codes, LLMs were consistently able to identify the general category of the clinical condition, although the specific subcode often differed from the dataset. Two-agent scenario is found to be a valuable resource for physicians, as it allows them to interact with an assistant that provides information and often suggests difficult diagnoses. It can be a useful tool in speeding up the diagnostic process.

5. Related Work

Prior studies explored multi-LLM frameworks to enhance accuracy and reasoning, primarily focusing on closed-ended questions [17, 18, 19, 20, 21, 22, 23, 24]. However, their applications remain confined to controlled settings, with limited exploration of real-world human-LLM collaboration. Evaluating LLMs’

multi-turn dialogue capabilities is a step toward practical applications. Kwan et al. [25] introduced the MT-Eval benchmark, finding that closed-source models outperform open-source ones, though multi-turn dialogues degrade performance due to retrieval difficulties and error propagation. Bai et al. [26] proposed MT-Bench-101 to assess LLMs in multi-turn dialogues, noting issues with adaptability and interactivity. Alignment techniques like RLHF [27] and DPO [28], as well as chat-specific designs, offered limited benefits for multi-turn tasks. Campedelli et al. [29] examined open-source LLMs in goal-driven collaborations and observed mixed success, with models like Mixtral [30] and Mistral [31] exhibiting higher failure rates.

In healthcare, LLMs have been explored for clinical note summarization [32, 33], aiming to assist physicians, though issues such as hallucinations and missing information persist [34, 35]. Additionally, metrics like ROUGE [36] and BLEU [37] used to assess summary quality have faced criticism regarding their effectiveness in evaluating clinical content. Furthermore, simulated patient-doctor interactions have been explored to enhance diagnostic accuracy. Liao et al. [38] improved accuracy by prompting LLMs to ask clarifying questions, though hallucinations persisted. Liu et al. [39] introduced the LLM-specific clinical pathway (LCP) to evaluate diagnostic performance using subjective and objective patient data, revealing challenges in handling multi-turn dialogues and clinical specialties, though their study focused solely on the Chinese language. Xie et al. [3] emphasized LLMs as supportive tools rather than replacements, developing the DoctorFLAN dataset and DotaBench to benchmark medical tasks. While most LLMs underperformed, DotaGPT, trained on DoctorFLAN, achieved superior results, demonstrating the dataset’s effectiveness. However, its availability only in Chinese limits the generalizability of the findings to other languages. Kim et al. [4, 5] proposed MDAgents, a framework that improves LLM effectiveness in complex medical decision-making by dynamically structuring collaboration models. It adapts to clinical needs by assigning LLMs independently or in groups based on task complexity. However, it fails to consider the critical role of physicians in medical decisions. Finally, Fan et al. [6] proposed the AI Hospital framework for simulated clinical diagnostics, whereas our approach focuses on iterative physician-LLM collaboration to refine clinical reasoning and decision-making.

6. Conclusion and Future Work

This work-in-progress paper introduced MedSyn, a dynamic human-AI collaboration framework designed to enhance clinical decision-making through multi-turn, conversational interactions between physicians and LLMs. Unlike traditional, static decision-support tools, MedSyn fosters an iterative diagnostic process where human expertise and AI-generated insights evolve together, aiming to create a safety net in complex medical scenarios. Through controlled simulations and qualitative analysis, we demonstrated that open-source LLMs are promising in meaningfully assisting physicians by uncovering overlooked information, proposing alternative hypotheses, and contributing to more comprehensive diagnostic reasoning. Our results revealed that while model performance varies, open-source LLMs show promise in improving diagnostic completeness and identifying rare conditions. In addition, physician evaluations highlighted the value of AI assistants not only in information retrieval, but also in hypothesis generation and diagnostic refinement. Despite encouraging results, challenges remain in aligning model outputs with clinical standards, particularly in the accurate generation of ICD-10 codes and managing nuances like chronic vs. acute conditions. These findings underscore the importance of continued iteration on evaluation metrics and dialogue strategies.

Future work will involve human-in-the-loop evaluations, enabling real physicians to engage with MedSyn in real-world settings and provide feedback on usability, relevance, and trustworthiness. We also plan to enhance MedSyn’s factual accuracy in clinical reasoning and coding, ensuring more robust and reliable support. This line of research is critical for the responsible integration of AI into clinical workflows—aiming to reduce diagnostic errors, support clinician decision-making, and ultimately improve patient outcomes. MedSyn represents a step toward more adaptive, intelligent healthcare systems where AI serves not as a replacement, but as a reliable and responsive partner in healthcare.

Acknowledgments

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HaDEA). Neither the European Union nor the granting authority can be held responsible for them. Grant Agreement no. 101120763 - TANGO. Andrea Passerini also acknowledges the support of the MUR PNRR project FAIR - Future AI Research (PE00000013) funded by the NextGenerationEU.

Declaration on Generative AI

During the preparation of this manuscript, the authors utilized ChatGPT and Grammarly to assist with paraphrasing, improving writing style, and refining grammar. After using these tools, the authors reviewed and edited the content as needed and took full responsibility for the publication's content.

References

- [1] G. Saposnik, D. Redelmeier, C. C. Ruff, P. N. Tobler, Cognitive biases associated with medical decisions: a systematic review, *BMC medical informatics and decision making* 16(1):138 (2016). doi:10.1186/s12911-016-0377-1.
- [2] A. N. Meyer, T. D. Giardina, L. Khawaja, H. Singh, Patient and clinician experiences of uncertainty in the diagnostic process: Current understanding and future directions, *Patient Education and Counseling* 104 (2021) 2606–2615. URL: <https://www.sciencedirect.com/science/article/pii/S0738399121004870>. doi:<https://doi.org/10.1016/j.pec.2021.07.028>.
- [3] W. Xie, Q. Xiao, Y. Zheng, X. Wang, J. Chen, K. Ji, A. Gao, X. Wan, F. Jiang, B. Wang, Llms for doctors: Leveraging medical llms to assist doctors, not replace them, *arXiv abs/2406.18034* (2024). URL: <https://arxiv.org/abs/2406.18034>.
- [4] Y. Kim, C. Park, H. Jeong, Y. S. Chan, X. Xu, D. McDuff, H. Lee, M. Ghassemi, C. Breazeal, H. W. Park, Mdagents: An adaptive collaboration of llms for medical decision-making, *arXiv abs/2404.15155* (2024). URL: <https://arxiv.org/abs/2404.15155>.
- [5] Y. Kim, C. Park, H. Jeong, C. Grau-Vilchez, Y. S. Chan, X. Xu, D. McDuff, H. Lee, C. Breazeal, H. W. Park, A demonstration of adaptive collaboration of large language models for medical decision-making, *arXiv abs/2411.00248* (2024). URL: <https://arxiv.org/abs/2411.00248>.
- [6] Z. Fan, L. Wei, J. Tang, W. Chen, W. Siyuan, Z. Wei, F. Huang, AI hospital: Benchmarking large language models in a multi-agent medical interaction simulator, in: O. Rambow, L. Wanner, M. Apidianaki, H. Al-Khalifa, B. D. Eugenio, S. Schockaert (Eds.), *Proceedings of the 31st International Conference on Computational Linguistics*, Association for Computational Linguistics, Abu Dhabi, UAE, 2025, pp. 10183–10213. URL: <https://aclanthology.org/2025.coling-main.680/>.
- [7] A. Johnson, L. Bulgarelli, L. Shen, A. Gayles, A. Shammout, S. Horng, T. Pollard, S. Hao, B. Moody, B. Gow, L.-w. Lehman, L. Celi, R. Mark, Mimic-iv, a freely accessible electronic health record dataset, *Scientific Data* 10 (2023) 1. doi:10.1038/s41597-022-01899-x.
- [8] A. Johnson, T. Pollard, S. Horng, L. A. Celi, R. Mark, Mimic-iv-note: Deidentified free-text clinical notes (version 2.2), *PhysioNet* (2023). doi:10.13026/1n74-ne17.
- [9] A. Goldberger, L. Amaral, L. Glass, S. Havlin, J. Hausdorg, P. Ivanov, R. Mark, J. Mietus, G. Moody, C.-K. Peng, H. Stanley, P. Physiobank, Components of a new research resource for complex physiologic signals, *PhysioNet* 101 (2000).
- [10] A. . M. Llama Team, The llama 3 herd of models, 2024. URL: <https://arxiv.org/abs/2407.21783>.
- [11] G. D. Gemma Team, Gemma 2: Improving open language models at a practical size., *arXiv abs/2501.12948* (2025).
- [12] D.-A. Team, Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning, *arXiv abs/2501.12948* (2025).

- [13] WHO, International Classification of Diseases (ICD), 2016. URL: <http://www.who.int/classifications/icd/en/>, accessed on 2021-04-14.
- [14] G. Wang, S. Cheng, X. Zhan, X. Li, S. Song, Y. Liu, Openchat: Advancing open-source language models with mixed-quality data, arXiv preprint arXiv:2309.11235 (2023).
- [15] Z. Chen, A. H. Cano, A. Romanou, A. Bonnet, K. Matoba, F. Salvi, M. Pagliardini, S. Fan, A. Köpf, A. Mohtashami, A. Sallinen, A. Sakhaeirad, V. Swamy, I. Krawczuk, D. Bayazit, A. Marmet, S. Montariol, M.-A. Hartley, M. Jaggi, A. Bosselut, Meditron-70b: Scaling medical pretraining for large language models, arXiv abs/2311.16079 (2023). URL: <https://arxiv.org/abs/2311.16079>.
- [16] J. Edin, A. Junge, J. D. Havtorn, L. Borgholt, M. Maistro, T. Ruotsalo, L. Maaløe, Automated medical coding on mimic-iii and mimic-iv: A critical review and replicability study, in: Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '23, Association for Computing Machinery, New York, NY, USA, 2023, p. 2572–2582. URL: <https://doi.org/10.1145/3539618.3591918>. doi:10.1145/3539618.3591918.
- [17] C.-M. Chan, W. Chen, Y. Su, J. Yu, W. Xue, S. Zhang, J. Fu, Z. Liu, Chateval: Towards better LLM-based evaluators through multi-agent debate, in: The Twelfth International Conference on Learning Representations, 2024. URL: <https://openreview.net/forum?id=FQepisCUWu>.
- [18] Y. Du, S. Li, A. Torralba, J. B. Tenenbaum, I. Mordatch, Improving factuality and reasoning in language models through multiagent debate, in: Proceedings of the 41st International Conference on Machine Learning, ICML'24, JMLR.org, 2024.
- [19] D. Jiang, X. Ren, B. Y. Lin, LLM-blender: Ensembling large language models with pairwise ranking and generative fusion, in: A. Rogers, J. Boyd-Graber, N. Okazaki (Eds.), Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Toronto, Canada, 2023, pp. 14165–14178. URL: <https://aclanthology.org/2023.acl-long.792/>. doi:10.18653/v1/2023.acl-long.792.
- [20] G. Li, H. A. Al Kader Hammoud, H. Itani, D. Khizbullin, B. Ghanem, Camel: communicative agents for "mind" exploration of large language model society, in: Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23, Curran Associates Inc., Red Hook, NY, USA, 2023.
- [21] T. Liang, Z. He, W. Jiao, X. Wang, Y. Wang, R. Wang, Y. Yang, S. Shi, Z. Tu, Encouraging divergent thinking in large language models through multi-agent debate, in: Y. Al-Onaizan, M. Bansal, Y.-N. Chen (Eds.), Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Miami, Florida, USA, 2024, pp. 17889–17904. URL: <https://aclanthology.org/2024.emnlp-main.992/>. doi:10.18653/v1/2024.emnlp-main.992.
- [22] Z. Liu, Y. Zhang, P. Li, Y. Liu, D. Yang, Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization, ArXiv abs/2310.02170 (2023). URL: <https://api.semanticscholar.org/CorpusID:276421095>.
- [23] Q. Sun, Z. Yin, X. Li, Z. Wu, X. Qiu, L. Kong, Corex: Pushing the boundaries of complex reasoning through multi-model collaboration, arXiv abs/2310.00280 (2023).
- [24] Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, L. Jiang, X. Zhang, S. Zhang, J. Liu, A. H. Awadallah, R. W. White, D. Burger, C. Wang, Autogen: Enabling next-gen LLM applications via multi-agent conversations, in: First Conference on Language Modeling, 2024. URL: <https://openreview.net/forum?id=BAakY1hNKS>.
- [25] W.-C. Kwan, X. Zeng, Y. Jiang, Y. Wang, L. Li, L. Shang, X. Jiang, Q. Liu, K.-F. Wong, MT-eval: A multi-turn capabilities evaluation benchmark for large language models, in: Y. Al-Onaizan, M. Bansal, Y.-N. Chen (Eds.), Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Miami, Florida, USA, 2024, pp. 20153–20177. URL: <https://aclanthology.org/2024.emnlp-main.1124/>. doi:10.18653/v1/2024.emnlp-main.1124.
- [26] G. Bai, J. Liu, X. Bu, Y. He, J. Liu, Z. Zhou, Z. Lin, W. Su, T. Ge, B. Zheng, W. Ouyang, MT-bench-101: A fine-grained benchmark for evaluating large language models in multi-turn dialogues, in: L.-W. Ku, A. Martins, V. Srikumar (Eds.), Proceedings of the 62nd Annual Meeting of the Association for

Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Bangkok, Thailand, 2024, pp. 7421–7454. URL: <https://aclanthology.org/2024.acl-long.401/>. doi:10.18653/v1/2024.acl-long.401.

- [27] T. Kaufmann, P. Weng, V. Bengs, E. Hüllermeier, A survey of reinforcement learning from human feedback abs/2312.14925 (2024). URL: <https://arxiv.org/abs/2312.14925>.
- [28] R. Rafailov, A. Sharma, E. Mitchell, S. Ermon, C. D. Manning, C. Finn, Direct preference optimization: your language model is secretly a reward model, in: Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23, Curran Associates Inc., Red Hook, NY, USA, 2023.
- [29] G. M. Campedelli, N. Penzo, M. Stefan, R. Dessì, M. Guerini, B. Lepri, J. Staiano, I want to break free! persuasion and anti-social behavior of llms in multi-agent settings with social hierarchy, arXiv abs/2410.07109 (2024). URL: <https://arxiv.org/abs/2410.07109>.
- [30] A. Q. Jiang, A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. de las Casas, E. B. Hanna, F. Bressand, G. Lengyel, G. Bour, G. Lample, L. R. Lavaud, L. Saulnier, M.-A. Lachaux, P. Stock, S. Subramanian, S. Yang, S. Antoniak, T. L. Scao, T. Gervet, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mixtral of experts, 2024. URL: <https://arxiv.org/abs/2401.04088>.
- [31] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, L. R. Lavaud, M.-A. Lachaux, P. Stock, T. L. Scao, T. Lavril, T. Wang, T. Lacroix, W. E. Sayed, Mistral 7b, arXiv abs/2310.06825 (2023). URL: <https://arxiv.org/abs/2310.06825>.
- [32] K. Krishna, S. Khosla, J. Bigham, Z. C. Lipton, Generating SOAP notes from doctor-patient conversations using modular summarization techniques, in: C. Zong, F. Xia, W. Li, R. Navigli (Eds.), Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Association for Computational Linguistics, Online, 2021, pp. 4958–4972. URL: <https://aclanthology.org/2021.acl-long.384/>. doi:10.18653/v1/2021.acl-long.384.
- [33] P. Cai, F. Liu, A. Bajracharya, J. Sills, A. Kapoor, W. Liu, D. Berlowitz, D. Levy, R. Pradhan, H. Yu, Generation of patient after-visit summaries to support physicians, in: N. Calzolari, C.-R. Huang, H. Kim, J. Pustejovsky, L. Wanner, K.-S. Choi, P.-M. Ryu, H.-H. Chen, L. Donatelli, H. Ji, S. Kurohashi, P. Paggio, N. Xue, S. Kim, Y. Hahm, Z. He, T. K. Lee, E. Santus, F. Bond, S.-H. Na (Eds.), Proceedings of the 29th International Conference on Computational Linguistics, International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 2022, pp. 6234–6247. URL: <https://aclanthology.org/2022.coling-1.544/>.
- [34] A. Ben Abacha, W.-w. Yim, Y. Fan, T. Lin, An empirical study of clinical note generation from doctor-patient encounters, in: A. Vlachos, I. Augenstein (Eds.), Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, Dubrovnik, Croatia, 2023, pp. 2291–2302. URL: <https://aclanthology.org/2023.eacl-main.168/>. doi:10.18653/v1/2023.eacl-main.168.
- [35] F. Moramarco, A. Papadopoulos Korfiatis, M. Perera, D. Juric, J. Flann, E. Reiter, A. Belz, A. Savkov, Human evaluation and correlation with automatic metrics in consultation note generation, in: S. Muresan, P. Nakov, A. Villavicencio (Eds.), Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Dublin, Ireland, 2022, pp. 5739–5754. URL: <https://aclanthology.org/2022.acl-long.394/>. doi:10.18653/v1/2022.acl-long.394.
- [36] C.-Y. Lin, ROUGE: A package for automatic evaluation of summaries, in: Text Summarization Branches Out, Association for Computational Linguistics, Barcelona, Spain, 2004, pp. 74–81. URL: <https://aclanthology.org/W04-1013/>.
- [37] K. Papineni, S. Roukos, T. Ward, W.-J. Zhu, Bleu: a method for automatic evaluation of machine translation, in: P. Isabelle, E. Charniak, D. Lin (Eds.), Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Philadelphia, Pennsylvania, USA, 2002, pp. 311–318. URL: <https://aclanthology.org/P02-1040/>. doi:10.3115/1073083.1073135.

- [38] Y. Liao, Y. Meng, H. Liu, Y. Wang, Y. Wang, An automatic evaluation framework for multi-turn medical consultations capabilities of large language models, arXiv abs/2309.02077 (2023). URL: <https://arxiv.org/abs/2309.02077>.
- [39] L. Liu, X. Yang, F. Li, C. Chi, Y. Shen, S. Lyu, M. Zhang, X. Ma, X. Lv, L. Ma, Z. Zhang, W. Xue, Y. Huang, J. Gu, Towards automatic evaluation for llms' clinical capabilities: Metric, data, and algorithm, in: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 5466–5475. URL: <https://doi.org/10.1145/3637528.3671575>. doi:10.1145/3637528.3671575.