

Multi-Agent Approach for Sepsis Management

Victor Iapascurta^{1,2}, Ion Fiodorov¹, Adrian Belii², Viorel Bostan¹

¹Department of Software Engineering and Automatics, Technical University of Moldova, Chisinnu, Republic of Moldova

²Department of Anesthesia and Intensive Care, Nicolae Testemitanu State University of Medicine and Pharmacy, Chisinnu, Republic of Moldova

Objectives: The high incidence of sepsis necessitates the development of practical decision-making tools for intensivists, especially during the early, critical phases of management. This study evaluates a multi-agent system intended to assist clinicians with antibiotic therapy and adherence to current sepsis management guidelines before diagnostic results become available. **Methods:** A multi-agent system incorporating three specialized agents was developed: a sepsis management agent, an antibiotic recommendation agent, and a sepsis guidelines compliance agent. A sepsis case from the MIMIC IV database, organized as a clinical vignette, was used to integrate and test these agents for generating management recommendations. The system leverages retrieval-augmented generation to improve decision-making through the integration of current literature and guidelines. **Results:** The application produced management recommendations for a sepsis case associated with pneumonia, including early initiation of broad-spectrum antibiotics and close monitoring for clinical deterioration. Two expert intensivists evaluated these recommendations as “acceptable” and reported moderate interrater agreement (Cohen’s kappa = 0.622, $p = 0.003$) across various aspects of recommendation usefulness. **Conclusions:** The multi-agent system shows promise in enhancing decision-making for sepsis management by optimizing antibiotic therapy and ensuring guideline compliance. However, reliance on a single case study limits the generalizability of the findings, highlighting the need for broader validation in diverse clinical settings to improve patient outcomes.

Keywords: Sepsis, Decision Support Systems, Clinical, Medical Informatics Applications, Critical Care, Artificial Intelligence

Submitted: January 25, 2025

Accepted: April 2, 2025

Corresponding Author

Victor Iapascurta

Department of Software Engineering and Automatics, Technical University of Moldova, 168, Stefan cel Mare bd, MD-2004, Chisinnu, Republic of Moldova. Tel: +373-69105713, E-mail: victor.iapascurta@doctorat.utm.md (<https://orcid.org/0000-0002-4540-7045>)

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<http://creativecommons.org/licenses/by-nc/4.0/>) which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

© 2025 The Korean Society of Medical Informatics

I. Introduction

Sepsis continues to be a leading cause of morbidity and mortality among critically ill patients, necessitating timely and effective management strategies [1,2]. Early identification and treatment are critical, yet challenges remain—especially when immediate diagnostic results are unavailable. This study presents a multi-agent system designed to support intensivists in making informed decisions about antibiotic therapy and in adhering to established guidelines during the initial phase of sepsis management.

Sepsis is a complex syndrome characterized by a dysregulated host response to infection that results in organ dysfunction [3,4]. Timely initiation of appropriate antimicrobial therapy is essential for improving patient outcomes, as delays significantly increase mortality rates. Recent studies indicate that the mortality rate in septic patients may increase by 7%

to 9% for every hour that appropriate antibiotic therapy is delayed [5,6]. Therefore, rapid decision-making is of critical importance.

In clinical practice, intensivists encounter challenges due to the overwhelming volume of information on sepsis management. Although guidelines provided by organizations such as the Surviving Sepsis Campaign offer valuable recommendations [7], adherence is often inconsistent. Moreover, the emergence of antimicrobial resistance complicates the selection of empirical therapy, necessitating a tailored approach based on local resistance patterns and patient-specific factors.

The aim of this study was to evaluate a multi-agent system intended to assist with antibiotic therapy and adherence to contemporary sepsis management guidelines. The system integrates three specialized agents—a sepsis management agent, an antibiotic recommendation agent, and a guidelines compliance agent. By employing retrieval-augmented generation (RAG) techniques [8,9], the system seeks to enhance the decision-making process and ensure that clinicians have access to the most current and relevant information.

II. Case Description

1. Clinical Vignette of the Case

We present the case of an 86-year-old woman with a history of chronic obstructive pulmonary disease who was admitted to the intensive care unit with severe pneumonia and suspected sepsis, as identified in the MIMIC IV database [10] (Case ID 10020786). The case has been organized as a clinical vignette (Table 1) for further analysis. Initial assessments revealed hypotension, critically low blood oxygen levels, and an elevated white blood cell count, prompting immediate intervention. A blood culture was obtained, although the results were pending during the decision-making process.

2. Multi-Agent System Architecture

Recently, multi-agent systems based on large language models (LLMs) have shown promising results in solving complex problems across various domains [11,12]. The multi-agent system developed in this study consists of three key agents:

- (1) Sepsis management agent: This agent evaluates the overall management strategy for sepsis cases by analyzing the clinical vignette and relevant literature.
- (2) Antibiotic recommendation agent: This agent analyzes patient data and retrieves pertinent literature to recommend appropriate antibiotic therapies in accordance with current guidelines.
- (3) Sepsis guidelines compliance agent: This agent reviews the proposed treatment to ensure it aligns with established sepsis management guidelines.

The architecture and data flow of the system are illustrated in Figure 1. Each agent interacts seamlessly to analyze the clinical vignette and generate management recommenda-

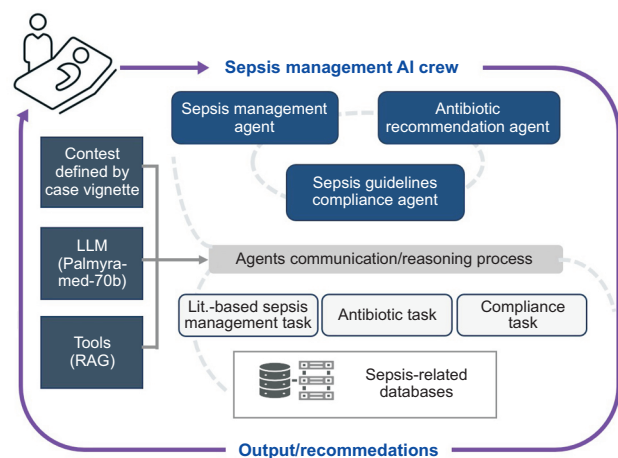


Figure 1. Architecture and data flow of the created multi-agent system. AI: artificial intelligence, LLM: large language model, RAG: retrieval-augmented generation.

Table 1. Clinical vignette of the sepsis case

Vignette section	Description
Diagnosis	Sepsis due to pneumonia
Demographic	Female, 86 years
Vital signs	Heart rate 76 bpm, Blood pressure 103/49 mmHg, Respiratory rate 22 rpm, Blood O ₂ saturation (SpO ₂) 90%, Temperature 37.2°C
Laboratory findings	White blood count 41,100/mm ³ (41.1 × 10 ⁹ cells/L), Blood glucose 225 mg/dL (12.5 mmol/L), Blood urea nitrogen 13 mg/dL (4.64 mmol/L)
Medical history	Chronic obstructive pulmonary disease, Atrial fibrillation, Type 2 diabetes mellitus without complications, Hypothyroidism
Special comments	Community-acquired pulmonary infection

tions.

The LLM is built on the Palmyra-med 70B large language model [13], and agent orchestration is facilitated by the CrewAI framework [14].

3. Retrieval-Augmented Generation

To increase the reliability of the recommendations, the agents employ a RAG approach, as depicted in Figure 2. This method leverages several persistent Chroma vector databases [15], which include recent literature and current sepsis management guidelines [16].

The system processes the clinical vignette along with a specific query, retrieving pertinent information from the sepsis-related databases to generate precise, evidence-based management recommendations tailored to the patient's condition.

4. Recommendations Generated by the System

Drawing on 20 recent, relevant articles related to sepsis that are organized within a persistent Chroma vector database [16], the system generated literature-based sepsis management recommendations. These recommendations include initiating early broad-spectrum antibiotics targeting likely pathogens in community-acquired pneumonia and closely monitoring the patient's clinical status. They also emphasize the importance of adjusting therapy based on culture results, continuous reassessment, and collaboration with a multidisciplinary team as central components of comprehensive care for this case. Figure 3A summarizes these recommendations.

Additionally, the system generated detailed antibiotic recommendations (Figure 3B) that emphasize the importance of tailoring empirical therapy for pulmonary infections by considering local microbial resistance patterns and patient-specific factors.

For the case of community-acquired pneumonia, the sys-

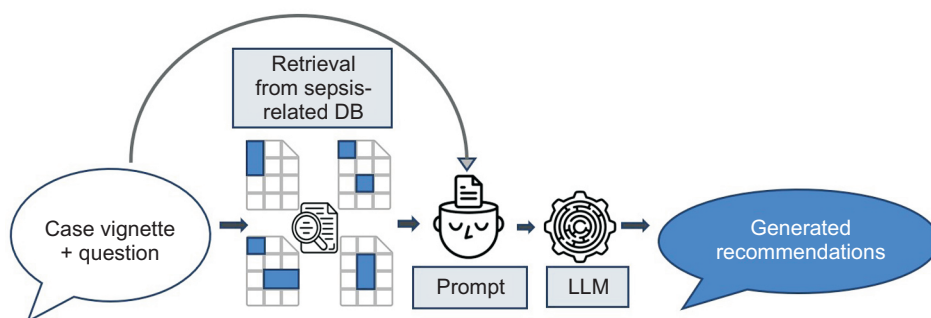


Figure 2. Retrieval-augmented generation (RAG) module. LLM: large language model.

A	Literature-based sepsis management recommendations
	Early administration of antibiotics is crucial in sepsis management, especially in patients with suspected infections and systemic inflammatory response syndrome criteria. In the case of sepsis due to pneumonia in an 86-year-old female with multiple comorbidities, including COPD and diabetes, prompt antibiotic therapy is essential. Considering the patient's condition, a broad-spectrum antibiotic covering typical pathogens in pneumonia should be initiated promptly. Close monitoring for clinical response and adjusting therapy based on culture results is recommended. Additionally, supportive care to stabilize vital signs and manage comorbidities is vital in improving outcomes for this patient. Regular reassessment and collaboration with a multidisciplinary team are key in providing comprehensive care for this complex case.
B	Antibiotic recommendations
	Tailoring empiric therapy for pulmonary infections, particularly Community-Acquired Pneumonia (CAP), should involve considering local microbial resistance patterns and individual patient factors. For CAP, initial empiric therapy may include a combination of a beta-lactam antibiotic like ceftriaxone or cefotaxime with a macrolide such as azithromycin. Alternatively, monotherapy with a respiratory fluoroquinolone like levofloxacin can be considered, especially for patients with beta-lactam allergies. In cases where there is a risk of MRSA or <i>Pseudomonas aeruginosa</i> , adding vancomycin or linezolid for MRSA and piperacillin-tazobactam, cefepime, or meropenem for <i>Pseudomonas</i> coverage may be necessary. It is crucial to adjust therapy based on local antimicrobial resistance data to ensure effective treatment.
C	Compliance with sepsis and septic shock management guidelines (2021)
	Prompt initiation of broad-spectrum antibiotics in sepsis is crucial, but de-escalation based on clinical improvement and culture results is recommended to prevent unnecessary therapy. Tailoring empiric therapy for pneumonia should consider local resistance patterns and patient factors to ensure effective treatment. To achieve compliance, consider the following suggestions: To enhance compliance with recommendations, ensure prompt initiation of broad-spectrum antibiotics for suspected pneumonia in the elderly patient with comorbidities. Tailor therapy based on local resistance patterns, consider beta-lactam/macrolide combination or fluoroquinolone for Community-Acquired Pneumonia. Adjust treatment according to culture results and collaborate with a multidisciplinary team for comprehensive care.

Figure 3. Recommendations generated by the multi-agent system: (A) sepsis management recommendations, (B) antibiotic recommendations, and (C) compliance with guidelines statement.




Figure 4. Evaluation results. LLM: large language model.

tem suggests a combination of beta-lactam antibiotics with a macrolide. Figure 3C outlines the recommendations for ensuring compliance with sepsis and septic shock management guidelines.

For convenience, a cloud-based application was developed using the agent definitions described in this paper [17]. Figure 3 illustrates the application's user interface and output. The application is available at <https://huggingface.co/Llm-RAGbasedAPPs> upon request to the corresponding author of this paper.

5. Evaluation Results

Since RAG plays a key role in generating the output for each agent, evaluation was performed using the TrueLens framework [18] based on the GPT-3.5-turbo LLM. The metric assessed answer relevance, context relevance (i.e., the usefulness of the context extracted from the vector store for generating the response), and groundedness (i.e., the extent to which the response is supported by the context). Figure 4 presents these evaluation results.

Additionally, two human experts evaluated the generated text using a similar metric. In this evaluation, context relevance and groundedness were combined into an overall answer-context groundedness score that reflects the alignment between the case description and the generated output.

III. Discussion

The multi-agent system successfully generated recommendations for a sepsis case attributed to pneumonia. Expert evaluation indicated that these recommendations were acceptable, suggesting the system's potential utility in real-world clinical settings. The observed moderate interrater

agreement among human experts (Cohen's kappa = 0.622, $p = 0.003$) suggests that the agent-generated outputs are generally consistent with expert judgment—a crucial factor in clinical decision-making. However, the formal agreement between the estimations provided by the LLM and those of the experts was negligible, warranting further investigation. Addressing this discrepancy may require careful selection of the materials used for constructing vector stores for RAG and exploring different settings during system development (e.g., chunk size, embedding model, similarity search function, and agent definitions).

1. Implications for Clinical Practice

The study findings indicate that the multi-agent system could serve as a valuable tool for intensivists, particularly in high-stakes scenarios where timely decision-making is critical. By offering rapid, evidence-based recommendations, the system has the potential to enhance the quality of care provided to patients with sepsis. In addition, its capacity to tailor antibiotic therapy based on local resistance patterns and patient-specific factors may lead to improved outcomes and lower rates of antimicrobial resistance.

2. Limitations and Future Directions

Nevertheless, this study has limitations. Reliance on a single case study restricts the generalizability of the findings to a broader patient population. Future research should focus on validating this approach across diverse clinical settings to fully realize its potential for enhancing sepsis management.

Furthermore, although the multi-agent system demonstrated promise in generating recommendations, its integration into clinical workflows must be carefully considered. Training and education for healthcare providers will be essential

to ensure effective utilization of the system and proper interpretation of its recommendations.

Additionally, traditional quantitative evaluation metrics may not fully capture the nuances of clinical decision-making, particularly in the context of sepsis.

Therefore, ground-truth evaluations by human experts, as demonstrated in this study, are recommended to assess the effectiveness of such decision-support systems.

3. Conclusion

In conclusion, the multi-agent system developed in this study holds the potential to improve decision-making in sepsis management. By integrating real-time data analysis with established guidelines, the system assists clinicians in delivering optimal care. However, further validation and integration into clinical practice are required to fully establish its efficacy.

By providing access to the code and documentation related to the system's development [16,17,19], the authors invite healthcare informatics researchers and clinicians to experiment with and enhance the application, ultimately benefiting patient care in sepsis management.

Conflict of Interest

No potential conflict of interest relevant to this article was reported.

Acknowledgments

This work was supported by the National Agency for Research and Development (NARD) of Moldova (Grant No. 25.00208.5007.07/PD) for the project "Intelligent Software Methods and Algorithms for Intensive Care: A Complexity Science Approach," implemented at the Technical University of Moldova.

ORCID

Victor Iapascurta (<https://orcid.org/0000-0002-4540-7045>)

Ion Fiodorov (<https://orcid.org/0000-0003-0938-3442>)

Adrian Belii (<https://orcid.org/0000-0002-4128-1318>)

Viorel Bostan (<https://orcid.org/0000-0002-2422-3538>)

References

1. Fleischmann C, Scherag A, Adhikari NK, Hartog CS, Tsaganos T, Schlattmann P, et al. Assessment of global incidence and mortality of hospital-treated sepsis. current estimates and limitations. *Am J Respir Crit Care Med* 2016;193(3):259-72. <https://doi.org/10.1164/rccm.201504-0781OC>
2. Strandberg G, Walther S, Agvald Ohman C, Lipcsey M. Mortality after severe sepsis and septic Shock in Swedish intensive care units 2008-2016: a nationwide observational study. *Acta Anaesthesiol Scand* 2020;64(7):967-75. <https://doi.org/10.1111/aas.13587>
3. Buckman SA, Turnbull IR, Mazuski JE. Empiric antibiotics for sepsis. *Surg Infect (Larchmt)* 2018;19(2):147-54. <https://doi.org/10.1089/sur.2017.282>
4. Singer M, Deutschman CS, Seymour CW, Shankar-Hari M, Annane D, Bauer M, et al. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *JAMA* 2016;315(8):801-10. <https://doi.org/10.1001/jama.2016.0287>
5. Liu VX, Fielding-Singh V, Greene JD, Baker JM, Iwashyna TJ, Bhattacharya J, et al. The timing of early antibiotics and hospital mortality in Sepsis. *Am J Respir Crit Care Med* 2017;196(7):856-63. <https://doi.org/10.1164/rccm.201609-1848OC>
6. Kumar A, Roberts D, Wood KE, Light B, Parrillo JE, Sharma S, et al. Duration of hypotension before initiation of effective antimicrobial therapy is the critical determinant of survival in human septic shock. *Crit Care Med* 2006;34(6):1589-96. <https://doi.org/10.1097/01.CCM.0000217961.75225.E9>
7. Evans L, Rhodes A, Alhazzani W, Antonelli M, Coopersmith CM, French C, et al. Surviving sepsis campaign: international guidelines for management of sepsis and septic shock 2021. *Intensive Care Med* 2021;47(11):1181-247. <https://doi.org/10.1007/s00134-021-06506-y>
8. Gao Y, Xiong Y, Gao X, Jia K, Pan J, Bi Y, et al. Retrieval-augmented generation for large language models: a survey [Internet]. Ithaca (CA): arXiv.org; 2024 [cited at 2025 Apr 13]. Available from: <https://doi.org/10.48550/arXiv.2312.10997>.
9. Iapascurta V, Fiodorov I. Retrieval-augmented generation using domain-specific text: a pilot study. *J Eng Sci* 2024;31(2):48-59. [https://doi.org/10.52326/jes.utm.2024.31\(2\).05](https://doi.org/10.52326/jes.utm.2024.31(2).05)
10. Johnson A, Bulgarelli L, Pollard T, Horng S, Celi LA, Mark R. (2023). MIMIC-IV Clinical Database Demo (version 2.2) [Internet]. York, UK: PhysioNet; 2023 [cited at 2025 Apr 13]. Available from: <https://doi.org/10.13026/dp1f-ex47>.
11. Guo T, Chen X, Wang Y, Chang R, Pei S, Chawla NV, et al. Large language model based multi-agents: A sur-

- vey of progress and challenges [Internet]. Ithaca (CA): arXiv.org; 2024 [cited at 2025 Apr 13]. Available from: <https://doi.org/10.48550/arXiv.2402.01680>.
12. Tang X, Zou A, Zhang Z, Li Z, Zhao Y, Zhang X, et al. Medagents: Large language models as collaborators for zero-shot medical reasoning [Internet]. Ithaca (CA): arXiv.org; 2024 [cited at 2025 Apr 13]. Available from: <https://doi.org/10.48550/arXiv.2311.10537>.
 13. AIModels. Palmyra-Med-70B [Internet]. AIModels; 2024 [cited at 2025 Apr 13]. Available from: <https://www.aimodels.fyi/models/huggingFace/palmyra-med-70b-writer>.
 14. CrewAI. The leading multi-agent platform [Internet]. San Francisco (CA): CrewAI; 2024 [cited at 2025 Jan 7]. Available from: <https://www.crewai.com>.
 15. Chroma. The AI-native open-source embedding database [Internet]. San Francisco (CA): Chroma; 2024 [cited at 2024 Dec 5]. Available from: <https://www.try-chroma.com>.
 16. Iapascurta V. Creating ChromaDBs sepsis-related [Internet]. San Francisco (CA): GitHub; 2024 [cited at 2025 Jan 7]. Available from: https://github.com/viapascurta-tech/Creating_ChromaDBs_sepsis_related.
 17. Iapascurta V. Palmyra med sepsis app [Internet]. San Francisco (CA): GitHub; 2024 [cited at 2025 Jan 7]. Available from: https://github.com/viapascurta-tech/Palmyra_med_sepsis_app.
 18. TruLens. Evaluate and track LLM applications [Internet]. TruLens; 2024 [cited at 2025 Jan 7]. Available from: <https://www.trulens.org>.
 19. Iapascurta V. Multi-agent approach [Internet]. San Francisco (CA): GitHub; 2024 [cited at 2025 Jan 7]. Available from: <https://github.com/viapascurta-tech/Multi-agent-approach>.