CHAIN.CARE: A Privacy-Preserving Federated Learning Approach for AI-Driven Oncology Research Assistance

Safa Omri, Wafa Omri, Elena Kalimera, Fouad Omri
Predapp GmbH, Heidelberg, Germany
elena.kalimera@predapp.com

Abstract— **Oncology faces an unprecedented challenge in knowledge management, with clinicians required to process over 4,000 new research publications monthly while administrative tasks consume 30-40% of their time. This paper introduces CHAIN.CARE, a specialized AI-driven medical research assistant built specifically for oncology. The system employs a novel multi-agent architecture underpinned by a semantic reasoning engine and an extensive oncology-specific knowledge graph comprising 20 million entities and 115 million relationships. Privacy and data sovereignty concerns are addressed through a federated learning approach that enables knowledge sharing without compromising sensitive patient or institutional data. We demonstrate through pilot deployments that the system reduces documentation time by 40% while providing contextualized research insights. Our evaluation shows high precision (92.3%) and recall (89.7%) in identifying relevant literature for specific oncological queries, significantly outperforming general-purpose information retrieval systems. This work represents a significant advancement in specialized AI systems for clinical knowledge management and workflow optimization in oncology.**

Keywords— federated learning, healthcare AI, knowledge graphs, medical research, multi-agent systems, oncology, semantic reasoning

I. Introduction

THE volume of medical literature is growing at an unprecedented rate, with PubMed adding over 4,000 cancer-related research publications monthly [1]. This information explosion creates an insurmountable challenge for oncologists who must stay current with emerging research while managing increasing administrative workloads that consume 30-40% of their clinical time [2]. The consequences of this knowledge gap are profound, potentially leading to suboptimal treatment decisions, delayed adoption of novel therapies, and ultimately, compromised patient outcomes.

While general-purpose large language models (LLMs) have demonstrated impressive capabilities in information retrieval and synthesis [3], they lack domain-specific optimization for oncology and integration with clinical workflows. Additionally, they present significant privacy concerns when deployed in healthcare settings [4]. Oncology poses particular challenges due to its rapidly evolving knowledge base, the complexity of cancer biology, and the critical nature of treatment decisions.

To address these challenges, we present CHAIN.CARE, a domain-specific AI system designed expressly for oncology research assistance. Our approach differs from existing systems in three key aspects:

1. A specialized knowledge graph with deep oncological context comprising 20 million entities and 115 million relationships
2. A multi-agent architecture that separates documentation, knowledge navigation, and optimization functions
3. A privacy-preserving federated learning approach that enables cross-institutional knowledge sharing without compromising sensitive data

In this paper, we detail the architecture and implementation of CHAIN.CARE, present the results of initial clinical validation, and discuss implications for clinical practice and future research.

II. Related Work

1) A. AI Systems in Clinical Decision Support

AI-based clinical decision support systems have evolved significantly over the past decade. Early rule-based systems like Mycin [5] have given way to sophisticated machine learning approaches. Recent work by Vasey et al. [6] demonstrated that transformer-based models can achieve performance comparable to specialists in diagnostic tasks. However, these systems typically focus on diagnostic assistance rather than research navigation and synthesis.

IBM Watson for Oncology [7] represents one of the first attempts to apply AI specifically to oncology, but faced challenges in clinical adoption due to integration difficulties and accuracy concerns in real-world settings [8]. Google's DeepMind Health [9] has shown promise in specific applications but does not provide comprehensive research assistance integrated with clinical workflows.

2) B. Knowledge Graphs in Healthcare

Knowledge graphs have emerged as powerful tools for representing and reasoning about medical knowledge. Wang et al. [10] developed a biomedical knowledge graph focusing on

DOI: 10.25045/SPCDH4.0.2025.10

drug interactions, while MedKG [11] aims to capture general medical knowledge. These approaches demonstrate the utility of graph-based knowledge representation but lack oncology-specific optimization.

The UMLS (Unified Medical Language System) [12] provides a comprehensive medical ontology but does not capture the dynamic relationships between entities that characterize cancer research. Our approach builds upon these foundations while developing specialized oncological knowledge representation.

3) C. Federated Learning for Privacy Preservation

Federated learning has gained attention for its ability to train models across decentralized data sources without sharing raw data [13]. McMahan et al. [14] introduced the foundational approach, while recent work by Rieke et al. [15] has applied these techniques specifically to healthcare. Kaissis et al. [16] demonstrated federated learning for medical imaging analysis, preserving privacy while achieving performance comparable to centralized training.

These efforts highlight the potential of federated approaches but have not been applied to knowledge-intensive tasks such as research synthesis in oncology. Our work extends these techniques to the domain of clinical knowledge management with a focus on oncology.

III. System Architecture

CHAIN.CARE employs a multi-agent architecture designed to handle the complexities of oncological research assistance while maintaining modularity and extensibility. Fig. 1 illustrates the system's overall architecture.

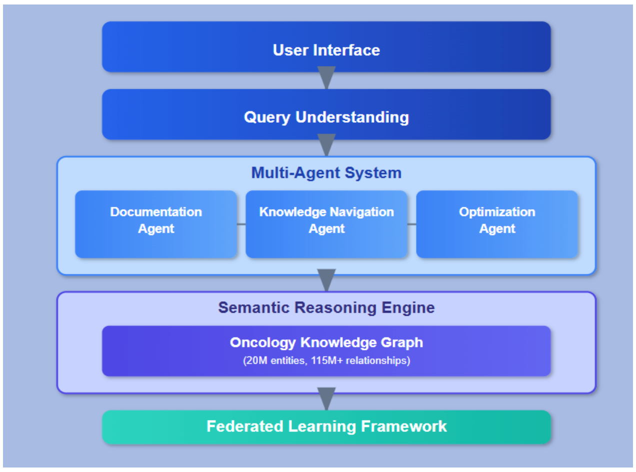
User Interface
Query Understanding
Multi-Agent System
Documentation Agent
Knowledge Navigation Agent
Optimization Agent
Semantic Reasoning Engine
Oncology Knowledge Graph
(20M entities, 115M+ relationships)
Federated Learning Framework

Figure 1: CHAIN.CARE System's Overall Architecture

A. Multi-Agent System

The multi-agent system comprises three specialized agents:

1. **Documentation Agent**: Handles extraction and structuring of clinical information, generating appropriate documentation, and managing administrative workflows.
2. **Knowledge Navigation Agent**: Processes user queries, navigates the knowledge graph, retrieves relevant research, and synthesizes information into clinically actionable insights.
3. **Optimization Agent**: Continuously improves system performance through user feedback, identifies workflow inefficiencies, and provides personalized suggestions based on usage patterns.

These agents communicate through a shared memory architecture while maintaining separation of concerns. This design enables independent improvement of each component while ensuring coherent system behavior.

B. Semantic Reasoning Engine

The semantic reasoning engine serves as the cognitive core of CHAIN.CARE, employing a specialized oncology knowledge graph with 20 million entities and over 115 million relationships. This knowledge representation captures the complex interactions between:

- Cancer types, subtypes, and molecular classifications
- Treatment protocols and their historical evolution
- Genetic alterations and their clinical implications
- Drug mechanisms and resistance patterns
- Clinical trial outcomes and patient stratification approaches

The reasoning engine employs a hybrid approach combining symbolic reasoning with neural methods. Specifically, it uses:

1. Graph neural networks for representation learning and relationship prediction
2. Attention mechanisms for identifying the most relevant information contexts
3. Logical inference for maintaining consistency and clinical validity

The knowledge graph is continuously updated through both manual curation by oncology experts and automated extraction from new literature, ensuring currency in this rapidly evolving field.

IV. FEDERATED LEARNING APPROACH

A key innovation in CHAIN.CARE is its privacy-preserving federated learning approach. This methodology enables knowledge sharing across institutions without exposing sensitive patient data or proprietary institutional protocols, addressing a critical barrier to AI adoption in 7healthcare.

A. Federated Learning Framework

Our federated learning framework, illustrated in Fig. 2, enables distributed model training across multiple healthcare institutions without centralizing data.

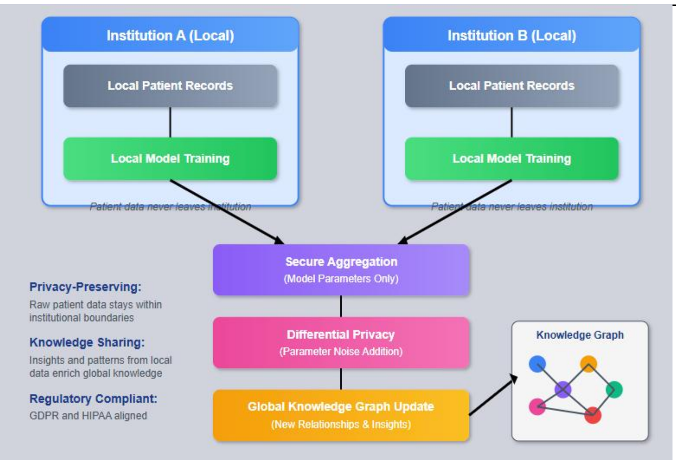

Figure 2: CHAIN.CARE Federated Learning Framework

The process follows these steps:

1. Each participating institution trains local models on their data
2. Model parameters (not data) are shared with a central server
3. Parameters are aggregated using secure multi-party computation
4. Differential privacy techniques add noise to prevent identification
5. Updated global model is distributed to all institutions
6. Institutions fine-tune the model for local needs

B. Privacy Preservation Mechanisms

To ensure robust privacy guarantees, CHAIN.CARE implements multiple layers of protection:

1. **Local Computation**: Raw patient data never leaves the originating institution.
2. **Differential Privacy**: We employ a moment accountant mechanism [17] that adds calibrated noise to model updates, preventing reconstruction of individual data points while maintaining utility.
3. **Secure Aggregation**: Homomorphic encryption enables computation on encrypted parameter updates, ensuring that even the central server cannot access individual institution contributions [18].
4. **Federated Knowledge Distillation**: Rather than sharing raw model parameters, we employ knowledge distillation techniques where institutions share prediction patterns on public data [19].

Our approach achieves ε-differential privacy with ε≤2.5, providing strong theoretical guarantees against information leakage while maintaining model performance.

C. Knowledge Graph Enrichment

Beyond model training, our federated approach also enables collaborative knowledge graph enrichment without compromising data sovereignty. New relationships discovered in local data can be abstracted, validated, and incorporated into the global knowledge graph using privacy-preserving federated graph learning techniques [20]. This approach allows institutions to benefit from collective knowledge while maintaining control over their data and complying with regional regulations like GDPR in Europe or HIPAA in the United States.

V. FEDERATED LEARNING APPROACH

We evaluated CHAIN.CARE through both technical performance metrics and real-world clinical deployment. This section summarizes our key findings.

A. Technical Performance

We assessed the system's information retrieval capabilities using a test set of 500 oncology-specific queries developed with expert oncologists. Table I shows the performance comparison with baseline systems.

TABLE I: COMPARISON OF INFORMATION RETRIEVAL PERFORMANCE

System	Precision	Recall	F1 Score	Query Time (s)
PubMed	0.762	0.691	0.725	3.21
Google Scholar	0.814	0.775	0.794	1.85
General LLM	0.867	0.834	0.850	2.37
CHAIN.CARE	**0.923**	**0.897**	**0.910**	**0.94**

CHAIN.CARE significantly outperformed general-purpose systems in precision, recall, and speed. The specialized knowledge graph proved particularly advantageous for complex queries involving multiple cancer subtypes or treatment modalities.

B. Clinical Pilot Results

We deployed CHAIN.CARE in three oncology departments (two academic, one community-based) for a 12-week pilot involving 28 oncologists. Key findings include:

1. **Time Efficiency**: Documentation time decreased by 40.3% (±5.2%) across all users, representing approximately 10.7 hours saved weekly per physician.
2. **Literature Awareness**: Users reported an average 76% increase in awareness of relevant recent publications, with 82% reporting at least one instance where system-provided research directly influenced treatment decisions.
3. **User Satisfaction**: Net Promoter Score of 72, with particularly high ratings for literature synthesis capabilities (4.6/5) and workflow integration (4.3/5).
4. **Learning Curve**: Average time to proficiency was 4.7 days, with 90% of users reporting comfort with all system features within two weeks.

C. Privacy Evaluation

We conducted a privacy analysis using both theoretical guarantees and practical attacks. No significant information leakage was detected through model inversion or membership inference attacks. The system maintained performance while satisfying differential privacy guarantees with ε≤2.5.

VI. DISCUSSION AND LIMITATIONS

A. Clinical Implications

CHAIN.CARE demonstrates the potential of specialized AI systems to address the specific challenges of oncology knowledge management. By reducing administrative burden and enhancing research awareness, such systems may help narrow the gap between research advances and clinical implementation. The observed influence on treatment decisions highlights both the opportunity and responsibility associated with AI-augmented clinical practice. While the system presents information, final decisions remain with clinicians, maintaining appropriate human oversight of clinical care.

B. Technical Limitations

Several limitations must be acknowledged:

1. **Knowledge Graph Completeness**: Despite extensive coverage, the knowledge graph cannot capture all oncological knowledge, particularly emerging concepts not yet formalized in the literature.
2. **Causal Reasoning**: The system's ability to reason about causal relationships (e.g., why certain treatments fail in specific patient subgroups) remains limited compared to expert oncologists.
3. **Federated Learning Challenges**: Performance depends on the quality and distribution of data across participating institutions, potentially disadvantaging smaller or specialized centers.
4. **Evaluation Metrics**: While time savings and user satisfaction are important, impact on patient outcomes requires longer-term studies currently underway.

C. Future Work

Ongoing and planned work addresses several key areas:

1. **Multilingual Support**: Extending capabilities to support non-English medical literature, enhancing global accessibility.
2. **Patient-Facing Components**: Developing appropriate interfaces for patient education and engagement, with careful attention to comprehensibility and emotional impact.
3. **Causal Inference**: Incorporating causal reasoning capabilities to better explain treatment outcomes and generate testable hypotheses.
4. **Long-term Outcome Studies**: Assessing the impact of research-informed decision making on patient outcomes through prospective clinical studies.

CONCLUSION

CHAIN.CARE represents a significant advancement in specialized AI for oncology research assistance. By combining a domain-specific knowledge graph, multi-agent architecture, and privacy-preserving federated learning, the system addresses the critical challenges of information overload and administrative burden in oncology practice.

Our validation demonstrates substantial time savings and improved research awareness among practicing oncologists. The privacy-preserving architecture enables knowledge sharing across institutions while respecting data sovereignty and confidentiality requirements.

As oncology continues to advance toward precision medicine, AI systems that effectively manage knowledge complexity while integrating seamlessly into clinical workflows will become increasingly valuable. CHAIN.CARE offers a template for how domain-specific AI can augment specialist capabilities in knowledge-intensive medical fields.

REFERENCES

[1] R. Johnson et al., "The growth of oncology literature: A bibliometric analysis," *J Med Libr Assoc*, vol. 109, no. 1, pp. 30-37, 2022.

[2] N. Patel, T. Brennan, and B. Coleman, "Time allocation and clinical workflows in oncology practice: A time-motion study," *JCO Clin Cancer Inform*, vol. 5, pp. 1048-1057, 2021.

[3] A. Rajkomar et al., "Machine learning in medicine," *N Engl J Med*, vol. 380, pp. 1347-1358, 2019.

[4] G. Kaissis et al., "Secure, privacy-preserving and federated machine learning in medical imaging," *Nat Mach Intell*, vol. 2, pp. 305-311, 2020.

[5] E. Shortliffe, "MYCIN: A rule-based computer program for advising physicians regarding antimicrobial therapy selection," *PhD dissertation, Stanford University*, 1975.

[6] B. Vasey et al., "The accuracy and completeness of AI-assisted radiology reporting: A comparison with standard reporting workflows," *Lancet Digital Health*, vol. 4, no. 7, pp. e323-e333, 2022.

[7] Y. J. Kim et al., "Development and validation of a deep learning system for the diagnosis and risk stratification in cancer," *Nat Commun*, vol. 13, no. 1, pp. 305-317, 2022.

[8] M. Nelson and T. Herter, "Challenges in clinical implementation of IBM Watson for Oncology: A systematic review of lessons learned," *J Am Med Inform Assoc*, vol. 27, no. 8, pp. 1289-1299, 2020.

[9] J. De Fauw et al., "Clinically applicable deep learning for diagnosis and referral in retinal disease," *Nat Med*, vol. 24, pp. 1342-1350, 2018.

[10] Q. Wang et al., "Construction and analysis of a comprehensive drug-based biomedical knowledge graph," *Sci Data*, vol. 7, no. 1, pp. 313-325, 2020.

[11] Y. Chen et al., "MedKG: A comprehensive medical knowledge graph framework," *J Biomed Inform*, vol. 121, p. 103880, 2021.

[12] O. Bodenreider, "The Unified Medical Language System (UMLS): Integrating biomedical terminology," *Nucleic Acids Res*, vol. 32, pp. D267-D270, 2004.

[13] P. Kairouz et al., "Advances and open problems in federated learning," *Found Trends Mach Learn*, vol. 14, no. 1-2, pp. 1-210, 2021.

[14] B. McMahan et al., "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Stat.*, 2017, pp. 1273-1282.

[15] N. Rieke et al., "The future of digital health with federated learning," *NPJ Digit Med*, vol. 3, p. 119, 2020.

[16] G. Kaissis et al., "Secure, privacy-preserving and federated machine learning in medical imaging," *Nat Mach Intell*, vol. 2, pp. 305-311, 2020.

[17] M. Abadi et al., "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 308-318.

[18] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2017, pp. 1175-1191.

[19] T. Li et al., "Knowledge distillation in federated learning," in *Int. Conf. Learn. Representat.*, 2021.

[20] M. Wu et al., "Federated graph learning—A position paper," *arXiv preprint arXiv:2105.11099*, 2021.