

# BUILDING A RAG SYSTEM FOR TISSUE ENGINEERING: INSIGHTS FROM DOMAIN-SPECIFIC TEXT AND SEPSIS MANAGEMENT

**Iapascurta Victor**<sup>1,2</sup>

<sup>1</sup> Department of Software Engineering and Automatics, Technical University, Chisinau, Republic of Moldova.

<sup>2</sup> Department of Anesthesia and Intensive Care, *Nicolae Testemițanu* State University of Medicine and Pharmacy, Chisinau, Republic of Moldova.

**Introduction:** Tissue engineering (TI) is a complex field that requires access to specific domain knowledge. Large Language Models (LLMs) are powerful but limited by their training data. Retrieval Augmented Generation (RAG) addresses this limitation by integrating LLMs with external knowledge sources. This work explores the design and implementation of a RAG system for TI, inspired by our successful applications in specific domains, including sepsis management. Such a system can greatly enhance the efficiency and accuracy of knowledge retrieval and generation in TI.

**Materials and Methods:** The RAG system for TI can be developed using a large language model augmented by a vector database of TI literature. The knowledge base is to be curated from research papers, textbooks, and other domain-specific texts. These texts can be processed to create embeddings stored in a vector database for efficient retrieval. When a user submits a query, the LLM generates an initial response and retrieves relevant information from the database to refine its output, ensuring accuracy and relevance. The RAG-based sepsis management system that can serve as an example that can be adapted to TI is available (after sending a request to [victor.iapascurta@usmf.md](mailto:victor.iapascurta@usmf.md)) at: [https://huggingface.co/spaces/LlmRAGbasedAPPs/LLM\\_MA\\_RAG\\_Sepsis](https://huggingface.co/spaces/LlmRAGbasedAPPs/LLM_MA_RAG_Sepsis).

**Results:** Our previous implementations of RAG in different domains have shown that integrating external knowledge significantly enhances the accuracy and relevance of LLM-generated responses. For TI, we anticipate this approach will similarly improve the system's ability to handle complex, domain-specific queries.

**Conclusions:** The development of a RAG system for tissue engineering represents a significant step forward in leveraging AI technologies for domain-specific applications. By combining the generative capabilities of LLMs with the precision of a curated knowledge base, we can create a tool that not only understands the language of TI but also possesses the depth of knowledge required to make meaningful contributions to the field. This work outlines the design, implementation, and expected benefits of such a system, drawing from our successful experiences in other domains.

**Keywords:** Tissue engineering, RAG, LLMs, Knowledge base, Vector database, Domain-specific knowledge, AI in biomedical engineering.