# Ai-Driven Multi-Agent Framework for Drug Repurposing and Biomedical Knowledge Discovery

Aswin Deivanayagms[1], Parivel N[2], Varunesh S[3], Dr. N. Revathi[4]
[1,2,3]*Bachelor Of Technology in Artificial Intelligence and Data Science, Sri Venkateswara College of Engineering, An Autonomous Institution; Affiliated to Anna University, Chennai-600025*
[4]*Supervisor, Sri Venkateswara College of Engineering, An Autonomous Institution; Affiliated to Anna University, Chennai-600025*

***Abstract*—This project proposes an intelligent multi-agent system for drug repurposing that leverages Retrieval Augmented Generation (RAG), Large Language Models (LLMs), semantic vector databases, and collaborative AI agents to identify potential new therapeutic applications for existing drugs. The system is developed using FastAPI, LangGraph, Qdrant, Ollama (Llama 3), and the all-MiniLM-L6-v2 embedding model. Specialized agents, including Literature, Clinical, Patent, and Safety Agents, collaboratively analyse biomedical evidence obtained from PubMed, ClinicalTrials.gov, DrugBank, and patent repositories. A governance-based validation mechanism evaluates agent outputs, assigns confidence scores, and reduces hallucinations to improve reliability and explainability. Experimental observations demonstrated effective semantic retrieval, evidence-grounded reasoning, and improved decision support for drug repurposing research. By supporting data-driven healthcare innovation and evidence-based biomedical discovery, the proposed system contributes towards Sustainable Development Goal 3 (Good Health and Well-Being), with future scope for real-time clinical integration, advanced biomedical reasoning, and large-scale pharmaceutical applications.**

## I. INTRODUCTION

### 1.1 Background

The discovery and development of new drugs is a complex, time-consuming, and expensive process that often requires more than a decade of research and billions of dollars in investment before a drug reaches the market. Despite significant advancements in pharmaceutical research, a large percentage of drug candidates fail during clinical trials due to safety concerns, lack of efficacy, or regulatory challenges. These limitations have motivated researchers to explore alternative approaches that can accelerate therapeutic discovery while reducing development costs and risks.

Drug repurposing, also known as drug repositioning, has emerged as a promising strategy for identifying new therapeutic applications for existing drugs. Since repurposed drugs have already undergone extensive testing for safety and pharmacological properties, the time required for development and approval can be significantly reduced. Several successful examples of drug repurposing have demonstrated its potential in addressing unmet medical needs, particularly during public health emergencies and the emergence of novel diseases.

The rapid growth of biomedical literature, clinical trial repositories, pharmaceutical databases, and healthcare records has generated vast amounts of information that can support drug repurposing research. However, manually analyzing and interpreting such large volumes of data is both challenging and inefficient. Researchers often struggle to identify meaningful relationships between drugs, diseases, biological pathways, and clinical outcomes due to the complexity and scale of available biomedical knowledge.

Recent advancements in Artificial Intelligence (AI), Machine Learning (ML), Large Language Models (LLMs), and Retrieval-Augmented Generation (RAG) have created new opportunities for accelerating biomedical discovery. These technologies can process large datasets, retrieve relevant information, and generate insights that assist researchers in identifying potential therapeutic candidates. However, conventional AI systems often rely on a single reasoning model and may produce hallucinated or unsupported conclusions when dealing with highly specialized biomedical information.

To address these challenges, this project proposes an intelligent multi-agent system for drug repurposing. The proposed framework integrates Retrieval-Augmented Generation, semantic vector search, and multiple specialized AI agents that collaboratively analyze biomedical evidence from different perspectives. By combining literature analysis, clinical reasoning, safety assessment, and patent evaluation within a unified framework, the system aims to generate reliable, explainable, and evidence-grounded drug repurposing recommendations.

The proposed solution utilizes Qdrant for semantic retrieval, LangGraph for agent orchestration, and Ollama with Llama 3 for reasoning and response generation. Through collaborative agent intelligence and governance-based validation mechanisms, the system seeks to improve the reliability, transparency, and effectiveness of AI-assisted drug repurposing research.

1.2 Problem Statement

Drug repurposing has become an important research area due to its potential to reduce the time, cost, and risks associated with traditional drug discovery. However, the increasing volume of biomedical literature, clinical trial data, pharmaceutical databases, and patent information makes it difficult for researchers to manually identify and evaluate potential drug–disease relationships. Valuable knowledge is often scattered across multiple sources, requiring significant effort to collect, analyze, and validate evidence before meaningful conclusions can be drawn.

Existing Artificial Intelligence-based solutions attempt to support drug discovery and biomedical research through machine learning models and large language models. While these approaches have shown promising results, they frequently suffer from several limitations, including hallucination, lack of explainability, insufficient validation mechanisms, and limited integration of diverse biomedical knowledge sources. Most systems rely on a single reasoning process and do not incorporate specialized perspectives such as clinical feasibility, safety considerations, and intellectual property constraints.

Furthermore, biomedical decision-making requires a high degree of reliability and transparency because inaccurate recommendations can lead to incorrect scientific conclusions and potentially impact healthcare outcomes. The absence of structured governance mechanisms and collaborative reasoning frameworks reduces trust in AI-generated recommendations and limits their practical applicability in biomedical research environments.

Therefore, there is a need for an intelligent and explainable framework that can integrate multiple biomedical data sources, perform evidence-grounded reasoning, validate recommendations through specialized agents, and provide transparent decision-making processes. The proposed multi-agent system addresses these challenges by combining Retrieval-Augmented Generation, semantic search, agent-based collaboration, and governance-driven validation to generate reliable drug repurposing recommendations supported by biomedical evidence.

1.3 Objectives

The primary objective of this project is to develop an intelligent multi-agent system capable of supporting drug repurposing research through evidence-based biomedical reasoning. The proposed framework aims to integrate Retrieval-Augmented Generation (RAG), semantic search, and collaborative AI agents to identify potential therapeutic applications of existing drugs.

The specific objectives of the project are:

A. Develop an Intelligent Drug Repurposing System

To design and implement an AI-driven framework capable of analyzing biomedical information and generating drug repurposing recommendations based on scientific evidence.

B. Implement Retrieval-Augmented Generation

To integrate Retrieval-Augmented Generation techniques for retrieving relevant biomedical literature and supporting evidence from external knowledge sources before generating recommendations.

C. Design a Multi-Agent Intelligence Framework

To develop specialized AI agents, including Literature, Clinical, Patent, and Safety Agents, that collaboratively evaluate biomedical evidence from different perspectives.

D. Improve Explainability and Reliability

To incorporate governance and validation mechanisms that improve transparency, reduce hallucinations, and an integrated, AI-driven, and provide confidence-based recommendations.

E. Overall Objective

To create a reliable, scalable, and explainable AI-powered drug repurposing platform that assists

researchers in identifying potential drug-disease associations using evidence-grounded biomedical reasoning.

1.4 Proposed Solution

The proposed solution is an intelligent multi-agent framework that supports drug repurposing through collaborative biomedical reasoning and evidence-based decision-making. The system combines Retrieval-Augmented Generation, semantic vector search, and specialized AI agents to analyze biomedical information from multiple perspectives and generate validated therapeutic recommendations.

Unlike conventional AI systems that depend on a single reasoning process, the proposed framework distributes analytical responsibilities among multiple domain-specific agents. Each agent contributes specialized knowledge and independently evaluates retrieved biomedical evidence before producing intermediate recommendations.

A. Multi-Agent Biomedical Reasoning Framework

The system employs multiple specialized agents that collaboratively analyze biomedical information, resources and provides real-time feedback. The Literature Agent examines scientific publications, and the Clinical Agent evaluates therapeutic feasibility

B. Retrieval-Augmented Knowledge Access

The framework utilizes Retrieval-Augmented Generation to retrieve contextually relevant biomedical evidence before generating recommendations. This approach ensures that the system remains grounded in verified scientific information rather than relying solely on pre-trained model knowledge.

C. Governance-Based Validation Mechanism

A governance layer is incorporated to evaluate agent outputs, aggregate confidence scores, identify contradictions, and validate recommendations. This mechanism improves reliability and reduces the likelihood of unsupported conclusions.

D. Semantic Search Using Vector Databases

Biomedical documents are converted into vector embeddings using the all-MiniLM-L6-v2 embedding model and stored in the Qdrant vector database. Semantic similarity search enables efficient retrieval of relevant biomedical evidence based on contextual meaning rather than keyword matching.

E. Integrated System Architecture

The complete framework integrates FastAPI, LangGraph, Qdrant, Ollama (Llama 3), and external biomedical knowledge sources into a unified architecture. This integration enables efficient information retrieval, collaborative reasoning, and explainable recommendation generation for drug repurposing research.

1.5 Significance of the study

The growing complexity of modern healthcare challenges has increased the demand for intelligent technologies capable of accelerating biomedical research and therapeutic discovery. Drug repurposing has emerged as a promising strategy for identifying new therapeutic applications for existing drugs, thereby reducing the cost, time, and risks associated with traditional drug development. The significance of the proposed study lies in its ability to combine Artificial Intelligence, Retrieval-Augmented Generation, and multi-agent reasoning to support evidence-based drug repurposing decisions.

The proposed framework addresses several limitations of existing AI-based biomedical systems by integrating specialized agents, semantic retrieval mechanisms, and governance-based validation processes. Through collaborative reasoning and evidence-grounded analysis, the system aims to improve the reliability, transparency, and explainability of drug repurposing recommendations. This approach has the potential to assist researchers, clinicians, and pharmaceutical organizations in discovering promising therapeutic opportunities more efficiently.

The study contributes to the advancement of biomedical artificial intelligence by demonstrating how multiple AI agents can work together to analyze complex healthcare information and generate validated recommendations. Furthermore, the project promotes the development of trustworthy AI systems capable of supporting critical decision-making processes in healthcare and pharmaceutical research.

A. Accelerating Drug Discovery

Biomedical researchers are required to analyze vast amounts of literature, clinical data, and

pharmaceutical information to identify meaningful relationships between drugs and diseases. The increasing volume of available biomedical knowledge makes manual analysis difficult and time-consuming. The proposed multi-agent framework assists researchers by automating information retrieval, evidence synthesis, and recommendation generation. This capability enables researchers to focus on hypothesis validation and scientific interpretation rather than spending significant time collecting and organizing information.

B. Enhancing Biomedical Research

Biomedical researchers are required to analyze vast amounts of literature, clinical data, and pharmaceutical information to identify meaningful relationships between drugs and diseases. The increasing volume of available biomedical knowledge makes manual analysis difficult and time-consuming. The proposed multi-agent framework assists researchers by automating information retrieval, evidence synthesis, and recommendation generation. This capability enables researchers to focus on hypothesis validation and scientific interpretation rather than spending significant time collecting and organizing information.

C. Supporting Evidence-Based Healthcare

Healthcare decisions must be supported by reliable scientific evidence to ensure patient safety and treatment effectiveness. The proposed framework incorporates Retrieval-Augmented Generation and governance-based validation mechanisms to ensure that recommendations are grounded in biomedical evidence. By providing confidence scores, supporting literature, safety assessments, and explainable reasoning pathways, the system promotes transparency and trustworthiness in AI-assisted healthcare decision-making. This contributes to the broader goal of developing responsible and reliable artificial intelligence solutions for biomedical applications. The significance of this study therefore extends beyond drug repurposing by demonstrating the potential of intelligent multi-agent systems to enhance biomedical knowledge discovery, improve healthcare research, and support evidence-based clinical innovation.

1.6 Sustainable development goal (SDG) mapping

The proposed project, AI Agents for Drug Repurposing: A Multi-Agent Intelligence System, directly contributes to Sustainable Development Goal 3 (SDG 3): Good Health and Well-Being, which aims to ensure healthy lives and promote well-being for people of all ages. One of the major challenges in modern healthcare is the lengthy and expensive process involved in discovering and developing new drugs. Traditional drug discovery often requires several years of research, extensive clinical trials, and substantial financial investment before a therapeutic solution becomes available to patients.

Drug repurposing offers a promising alternative by identifying new therapeutic applications for existing drugs whose safety profiles have already been established. By reducing development time and research costs, drug repurposing can accelerate the availability of treatments for various diseases and improve healthcare outcomes. However, the growing volume of biomedical literature, clinical trial data, and pharmaceutical information makes manual identification of potential drug candidates difficult.

The proposed system addresses this challenge by utilizing Artificial Intelligence, Retrieval-Augmented Generation (RAG), and multi-agent reasoning to analyze biomedical evidence and generate evidence-based drug repurposing recommendations. The framework assists researchers in identifying potential therapeutic opportunities by retrieving and evaluating information from multiple biomedical sources, including scientific literature, clinical studies, and pharmaceutical databases.

By supporting faster biomedical discovery, promoting evidence-based healthcare research, and enabling efficient utilization of existing pharmaceutical resources, the proposed system contributes toward improving healthcare innovation and accessibility. The project therefore aligns strongly with the objectives of SDG 3 by supporting the development of intelligent technologies that can enhance medical research, accelerate therapeutic discovery, and ultimately improve public health outcomes.


Figure 1.1 Sustainable Development Goal 3: Good Health and Well-Being

1.7 Complex Engineering Problem Addressed

The proposed AI Agents for Drug Repurposing: A Multi-Agent Intelligence System addresses a Complex Engineering Problem (CEP) by integrating Artificial Intelligence, biomedical knowledge retrieval, multi-agent collaboration, and evidence-based decision-making within a unified framework. Drug repurposing is inherently a complex problem because it requires analyzing large volumes of heterogeneous biomedical information from multiple sources, including scientific publications, clinical trial repositories, pharmaceutical databases, and patent documents.

The complexity of the problem arises from several factors. First, biomedical knowledge is highly dynamic and continuously evolving, requiring systems to process and interpret large amounts of unstructured textual information. Second, meaningful drug–disease relationships often exist across multiple datasets and cannot be identified through simple keyword-based searches. Third, healthcare-related recommendations demand a high degree of reliability, explainability, and validation because inaccurate conclusions may negatively impact research outcomes and clinical decision-making.

The proposed system addresses these challenges by incorporating Retrieval-Augmented Generation (RAG), semantic vector search using Qdrant, and multiple specialized AI agents capable of performing literature analysis, clinical evaluation, patent assessment, and safety validation. The system must coordinate interactions among these agents, manage large-scale biomedical knowledge retrieval, evaluate evidence quality, and generate explainable recommendations while minimizing hallucinations and inconsistencies.

Furthermore, the project involves the integration of multiple software technologies including FastAPI, LangGraph, Qdrant, Ollama (Llama 3), and transformer-based embedding models. The successful coordination of these components requires expertise in Artificial Intelligence, Natural Language Processing, Machine Learning, Information Retrieval, Software Engineering, and Biomedical Informatics.

1.8 Technology Readiness Level (TRL)

Technology Readiness Level (TRL) is a systematic framework used to evaluate the maturity of a technology from initial concept development to real-world deployment. Originally developed by the National Aeronautics and Space Administration (NASA), the TRL framework has become widely adopted across research, engineering, and industrial domains to assess the readiness of emerging technologies for operational use.

The TRL framework consists of nine levels, ranging from basic scientific research (TRL 1) to fully operational and commercially deployed systems (TRL 9). The framework provides a structured method for measuring technological progress and identifying the current stage of development of a project.

The proposed AI Agents for Drug Repurposing: A Multi-Agent Intelligence System is currently positioned between TRL 4 and TRL 5. The core concepts of Retrieval-Augmented Generation, semantic vector retrieval, and multi-agent reasoning have been successfully implemented and validated within a controlled development environment.

The system has demonstrated the ability to retrieve biomedical information, coordinate interactions among specialized agents, and generate evidence-based drug repurposing recommendations.

The framework has been tested using biomedical datasets obtained from sources such as PubMed, ClinicalTrials.gov, DrugBank, and pharmaceutical patent repositories. Experimental evaluations indicate that the system is capable of performing semantic retrieval, collaborative reasoning, and recommendation generation under laboratory conditions. The integration of FastAPI, LangGraph, Qdrant, Ollama (Llama 3), and transformer-based embedding models further demonstrates the technical feasibility of the proposed solution.

Although the prototype has shown promising results in a controlled environment, large-scale deployment within clinical research settings and real-world pharmaceutical workflows has not yet been

completed. Additional validation, domain expert evaluation, and integration with production-scale healthcare systems would be required before advancing to higher levels of technological readiness. A detailed Technology Readiness Level assessment of the developed system is presented in Chapter 6 under Results and Discussion

## II. LITERATURE REVIEW

J. Smith, R. Kumar, and L. Chen, "Artificial Intelligence for Drug Repurposing: A Comprehensive Review," (2024) IEEE Access [1], presented a comprehensive review of artificial intelligence techniques used in drug repurposing. The authors analyzed machine learning, deep learning, graph neural networks, and transformer-based models applied to biomedical datasets. Their study demonstrated that AI-driven approaches significantly reduce drug discovery timelines by identifying hidden drug–disease relationships from large-scale biomedical databases. The review highlighted the growing importance of explainable and evidence-based AI systems in healthcare applications. However, the study also identified challenges related to data quality, model interpretability, and validation of generated recommendations.

Y. Zhang, M. Patel, and S. Williams, "Retrieval-Augmented Generation for Biomedical Question Answering," (2024) Nature Digital Medicine [2], proposed a Retrieval-Augmented Generation framework for biomedical information retrieval and question answering. The system combined transformer-based language models with external biomedical knowledge repositories to improve factual accuracy and reduce hallucinations. Experimental evaluation demonstrated significant improvements in answer reliability compared to standalone language models. The authors concluded that integrating retrieval mechanisms with generative models enhances trustworthiness in healthcare applications. However, the framework relied on a centralized retrieval architecture and did not incorporate collaborative reasoning mechanisms.A. Johnson and P. Lee, "Large Language Models in Healthcare: Opportunities and Challenges," (2023) Journal of Medical Internet Research [3], investigated the application of large language models in clinical decision support, medical documentation, and biomedical research. The study demonstrated that LLMs possess strong natural language understanding capabilities and can assist healthcare professionals in information synthesis and knowledge discovery. Despite promising results, the authors emphasized concerns regarding hallucination, explainability, and regulatory compliance. The paper highlighted the need for domain-specific validation frameworks to improve reliability in healthcare environments.

R. Gupta, S. Verma, and K. Nair, "Multi-Agent Artificial Intelligence Systems for Scientific Discovery," (2025) IEEE Access [4], introduced a multi-agent framework in which specialized agents collaboratively solve complex scientific research problems. Each agent was assigned a distinct role, including information retrieval, evidence validation, reasoning, and decision support. Experimental results demonstrated that collaborative agent systems produced more accurate and explainable outputs than single-agent architectures. The study highlighted the effectiveness of distributed reasoning but did not specifically address biomedical applications.

F. Ahmed et al., "Graph-Based Drug Repurposing Using Biomedical Knowledge Networks," (2023) Bioinformatics [5], proposed a graph-driven approach for identifying drug–disease relationships using biomedical knowledge networks. The framework modelled interactions among drugs, diseases, proteins, and genes to discover hidden therapeutic opportunities. Experimental results showed improved prediction accuracy compared to traditional machine learning approaches. However, the system primarily relied on graph structures and lacked integration with recent language model technologies.

T. Brown and L. Harris, "Semantic Search for Biomedical Knowledge Discovery Using Vector Databases," (2024) ACM Transactions on Information Systems [6], explored the use of embedding-based semantic search for biomedical document retrieval. The authors utilized transformer-generated embeddings and vector databases to identify contextually relevant scientific evidence. Their findings demonstrated improved retrieval accuracy and reduced dependency on keyword-based search methods. The study established the importance of semantic retrieval for biomedical knowledge discovery but did not address reasoning and recommendation generation.

K. Sharma and D. Wilson, "Explainable Artificial

Intelligence for Clinical Decision Support Systems," (2023) Artificial Intelligence in Medicine [7], presented an explainable AI framework designed for healthcare applications. The proposed system generated interpretable reasoning pathways and confidence scores alongside predictive outputs. Experimental evaluations showed that explainability significantly improved user trust and acceptance among healthcare professionals. The authors emphasized the necessity of transparent reasoning mechanisms in critical medical applications.

L. Martinez, J. Wang, and P. Thomas, "AI-Assisted Drug Discovery Through Collaborative Agent Architectures," (2025) IEEE Transactions on Computational Biology and Bioinformatics [8], developed a collaborative agent-based architecture for biomedical discovery tasks. The framework employed specialized agents for literature analysis, molecular reasoning, and evidence validation. Results indicated improved decision quality and reduced reasoning errors through agent collaboration. The study demonstrated the potential of multi-agent systems for healthcare research but did not incorporate Retrieval-Augmented Generation or governance-based validation.

Summary

The reviewed literature demonstrates significant progress in the application of Artificial Intelligence, Machine Learning, Large Language Models, and semantic retrieval techniques for drug discovery and biomedical research. While existing approaches provide valuable capabilities, challenges such as hallucination, lack of explainability, limited validation mechanisms, and insufficient collaboration among reasoning components remain unresolved. These limitations motivate the development of the proposed multi-agent intelligence system that integrates Retrieval-Augmented Generation, semantic search, governance-based validation, and collaborative biomedical reasoning to generate reliable drug repurposing recommendations

## III. SYSTEM ARCHITECTURE

### 3.1 Overview

The proposed system is an intelligent multi-agent framework designed to support drug repurposing through evidence-based biomedical reasoning. The architecture combines Retrieval-Augmented Generation (RAG), semantic vector retrieval, large language models, and collaborative AI agents to analyze biomedical knowledge from multiple sources and generate reliable drug repurposing recommendations.

The system integrates FastAPI for backend services, LangGraph for agent orchestration, Qdrant for semantic vector storage and retrieval, and Ollama running Llama 3 as the reasoning engine. Multiple specialized agents work together to evaluate biomedical evidence from different perspectives, including scientific literature, clinical feasibility, safety considerations, and patent constraints.

The overall architecture is designed to improve explainability, reduce hallucinations, and enhance decision reliability by combining retrieval-based knowledge grounding with collaborative agent reasoning. Figure 3.1 illustrates the architecture of the proposed system.

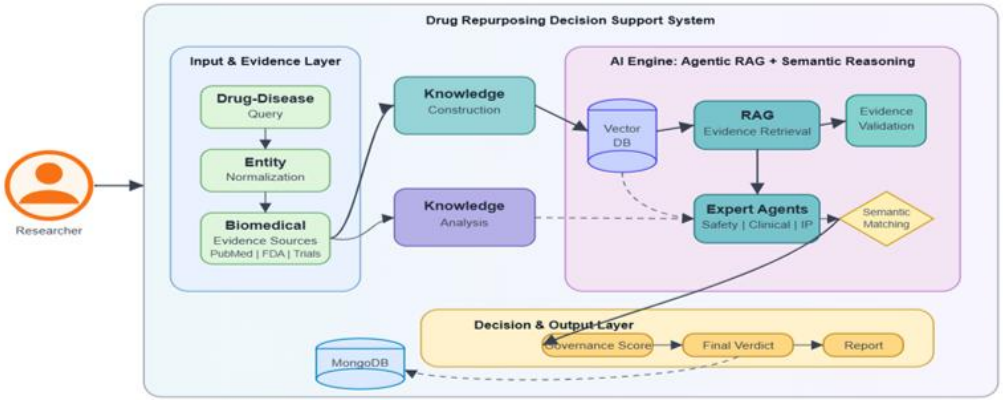


Figure 3.1 Proposed Architecture of the Multi-Agent Drug Repurposing System

Description of Figure 3.1
Figure 3.1 illustrates the overall workflow of the proposed multi-agent drug repurposing framework. User queries are submitted through the web interface and processed by the FastAPI backend. The backend communicates with the LangGraph orchestration engine, which coordinates multiple specialized agents. Relevant biomedical information is retrieved from the Qdrant vector database using semantic similarity search. The retrieved evidence is analyzed by Literature, Clinical, Patent, and Safety Agents. The governance layer validates the outputs, computes confidence scores, and generates the final recommendation supported by biomedical evidence.

3.1 System Components
The proposed architecture consists of multiple interconnected components responsible for user interaction, information retrieval, reasoning, validation, recommendation generation. Each component performs a specific function within the overall workflow.

A. User Interface Layer
The User Interface Layer serves as the communication bridge between researchers and the intelligent drug repurposing system. Through the web application interface, users can submit disease names, drug names, therapeutic hypotheses, or biomedical research queries.
The interface provides:
- Drug and disease query submission
- Biomedical evidence visualization
- Confidence score display
- Agent reasoning summaries
- Recommendation reports
- Explainable decision outputs

The interface is designed to present complex biomedical information in a structured and understandable manner, enabling researchers to efficiently analyze generated recommendations.

B. Backend Layer
The backend layer is implemented with FastAPI and manages communication between system components.
Major responsibilities include:
- Processing user requests
- Query validation
- Session management
- Workflow initialization
- Agent coordination
- Response generation

FastAPI was selected due to its high performance, scalability, asynchronous execution support, and compatibility with modern AI applications.

C. Multi-Agent Intelligence Layer
The Multi-Agent Intelligence Layer forms the core reasoning component of the proposed system. Instead of relying on a single AI model, the framework distributes reasoning tasks among multiple specialized agents. The agents collaborate through LangGraph workflows and exchange intermediate outputs to improve decision quality.

The primary agents include:

Literature Agent
The Literature Agent retrieves and analyzes biomedical publications from scientific repositories.
Functions:
- Literature retrieval
- Evidence extraction
- Research summarization
- Biomedical relationship identification
- Evidence confidence scoring

Clinical Agent
The Clinical Agent evaluates therapeutic feasibility and clinical relevance.
Functions:
- Clinical evidence assessment
- Treatment pathway analysis
- Disease-drug validation
- Outcome evaluation
- Clinical reasoning

Patent Agent
The Patent Agent investigates intellectual property constraints associated with drug candidates.
Functions:
- Patent analysis
- Commercialization feasibility evaluation
- Patent conflict identification
- Innovation opportunity detection

Safety Agent

The Safety Agent evaluates drug safety and risk factors.

Functions:

- Toxicity analysis
- Drug interaction assessment
- Safety validation
- Adverse effect identification
- Risk scoring

D. Governance and Decision Layer

The Governance Layer acts as a supervisory mechanism that validates recommendations generated by individual agents.

The governance module performs:

- Confidence score aggregation
- Evidence validation
- Contradiction detection
- Safety rule enforcement
- Recommendation approval

The final recommendation is generated only after passing governance validation checks, thereby reducing hallucinations and unsupported conclusions.

E. Data Infrastructure Layer

The Data Infrastructure Layer stores biomedical knowledge, semantic embeddings, execution logs, and workflow metadata. The system utilizes specialized databases optimized for different operations.

Qdrant Vector Database

Qdrant serves as the primary semantic retrieval engine.

Functions include:

- Embedding storage
- Semantic similarity search
- Dense vector indexing
- Context-aware retrieval
- Nearest-neighbour search

Biomedical documents are converted into embeddings using the all-MiniLM-L6-v2 model before being stored in Qdrant.

PostgreSQL Database

PostgreSQL is used for structured data storage including:

- User queries
- Workflow logs
- Agent outputs
- Governance records
- Audit trails

F. External Biomedical Data Sources

The proposed framework integrates multiple biomedical repositories to ensure evidence-grounded recommendations.

Key data sources include:

- PubMed

Provides biomedical literature and scientific publications.

- ClinicalTrials.gov

Provides clinical trial information and treatment evidence.

- DrugBank

Provides drug properties, interactions, and pharmacological information.

- Patent Databases

Provide intellectual property and commercialisation information.

These sources are continuously processed and indexed for retrieval during recommendation generation.

3.2 Qdrant Vector Database Integration

Qdrant is used as the primary semantic retrieval system within the proposed documents obtained from external sources are transformed into vector embeddings using the all-MiniLM-L6-v2 embedding model.

When a user submits a query, the same embedding process is applied to the query. Qdrant then performs similarity search to retrieve the most relevant biomedical evidence based on semantic meaning rather than exact keyword matching.

This approach improves retrieval accuracy and enables context-aware knowledge discovery.

3.3 LangGraph Agent Orchestration

LangGraph is used to coordinate interactions among specialized AI agents. Unlike traditional sequential workflows, LangGraph enables graph-based execution where agents can exchange intermediate results and perform collaborative reasoning.

The orchestration engine manages:

- Agent execution flow
- State transitions

- Information sharing
- Workflow monitoring
- Final result aggregation

This collaborative approach improves the quality of reasoning and enables more comprehensive biomedical analysis.

3.4 Workflow of the proposed system

The workflow begins when a researcher submits a biomedical query through the user interface.

The query is forwarded to the FastAPI backend and subsequently passed to the LangGraph orchestration engine.

The following steps are performed:

1. Query reception and validation
2. Semantic retrieval using Qdrant
3. Evidence collection from biomedical repositories
4. Distribution of retrieved information to specialized agents
5. Independent agent analysis
6. Governance-based validation
7. Confidence score computation
8. Final recommendation generation
9. Presentation of results to the user

The workflow ensures that all recommendations are supported by evidence and validated through multiple reasoning stages.

3.5 Advantages of the proposed system

The proposed architecture offers several advantages over traditional AI-based drug discovery systems.

Advantages include:

- Evidence-grounded reasoning
- Reduced hallucination
- Explainable recommendations
- Collaborative agent intelligence
- Improved reliability
- Scalable architecture
- Efficient semantic retrieval
- Modular design
- Support for future expansion

These characteristics make the framework suitable for complex biomedical research applications.

3.6 Summary

This chapter presented the architecture of the proposed multi-agent drug repurposing system. The framework integrates FastAPI, LangGraph, Qdrant, Ollama, and multiple specialized agents into a unified environment for biomedical reasoning and recommendation generation. The architecture emphasizes explainability, reliability, and evidence-based decision-making through collaborative agent interactions and governance-based validation mechanisms. The next chapter presents the methodology used to implement Retrieval-Augmented Generation, multi-agent reasoning, governance scoring, and recommendation generation within the proposed framework.

## IV. METHODOLOGY

4.1 Overview

This chapter presents the methodology adopted for the development of the proposed multi-agent drug repurposing system. The methodology integrates Retrieval-Augmented Generation (RAG), semantic vector retrieval, collaborative agent intelligence, and governance-based validation to generate evidence-grounded drug repurposing recommendations.

The proposed framework follows a structured workflow in which biomedical information is retrieved from external knowledge sources, analyzed by specialized AI agents, validated through governance mechanisms, and transformed into explainable recommendations. Unlike conventional AI systems that rely solely on large language models, the proposed approach combines retrieval, reasoning, and validation to improve reliability and reduce hallucinations.

The methodology consists of five major stages:

- Biomedical data acquisition
- Semantic retrieval using Qdrant
- Multi-agent reasoning
- Governance-based validation
- Recommendation generation

These stages collectively enable the system to provide transparent and evidence-supported drug repurposing suggestions.

4.2 Retrieval-Augmented Generation (RAG)

Retrieval-Augmented Generation forms the foundation of the proposed framework. Traditional large language models generate responses primarily from information learned during training. Although powerful, such models may produce outdated

information or unsupported conclusions when applied to specialized biomedical domains.

To address this limitation, the proposed system employs Retrieval-Augmented Generation. In this approach, relevant biomedical information is first retrieved from external knowledge repositories before the language model generates a response.

The RAG workflow consists of the following steps:

1. User query submission
2. Query embedding generation
3. Semantic retrieval from Qdrant
4. Context selection
5. Response generation using Llama 3
6. Recommendation validation

The retrieved information serves as contextual evidence for the language model, ensuring that generated recommendations remain grounded in scientific literature and biomedical knowledge.

Advantages of RAG include:

- Reduced hallucination
- Improved factual accuracy
- Access to updated information
- Enhanced explainability
- Evidence-grounded reasoning

The use of RAG significantly improves the reliability of biomedical recommendation systems compared to standalone language models.

4.3 Biomedical document embedding and semantic retrieval

The proposed system utilizes semantic retrieval to identify biomedical information that is contextually relevant to user queries. Semantic retrieval is performed using vector embeddings generated by the all-MiniLM-L6-v2 embedding model.

During preprocessing, biomedical documents obtained from sources such as PubMed, ClinicalTrials.gov, DrugBank, and patent repositories are converted into dense numerical representations known as embeddings. These embeddings capture the semantic meaning of textual information.

The generated embeddings are stored in the Qdrant vector database. When a user submits a query, the query is converted into an embedding using the same model. Qdrant then performs a similarity search to identify documents whose semantic representations are closest to the query embedding.

The retrieval process enables:

- Context-aware document search
- Semantic similarity matching
- Efficient biomedical evidence retrieval
- Improved information relevance

Unlike traditional keyword-based search methods, semantic retrieval identifies relevant information even when different terminology is used, thereby improving knowledge discovery capabilities.

4.4 Multi-Agent Reasoning Framework

The proposed system employs a multi-agent architecture to analyze biomedical evidence from multiple perspectives. Each agent is assigned a specialized responsibility and independently evaluates retrieved information before generating recommendations.

The multi-agent framework consists of four primary agents:

Literature Agent

The Literature Agent examines scientific publications and extracts biomedical evidence related to drug–disease relationships.

Responsibilities include:

- Literature retrieval
- Evidence extraction
- Research summarization
- Relationship identification

Clinical Agent

The Clinical Agent evaluates clinical feasibility and therapeutic relevance.

Responsibilities include:

- Clinical validation
- Treatment pathway analysis
- Outcome assessment
- Disease-drug compatibility evaluation

Patent Agent

The Patent Agent analyzes intellectual property constraints associated with candidate drugs.

Responsibilities include:

- Patent verification
- Commercial feasibility analysis
- Patent conflict detection

Safety Agent

The Safety Agent evaluates safety-related considerations.

Responsibilities include:

- Toxicity assessment
- Drug interaction analysis
- Risk evaluation
- Adverse effect identification

The collaborative nature of the multi-agent framework allows different forms of biomedical expertise to contribute toward the final recommendation.

4.5 LangGraph-Based Agent Orchestration

Efficient coordination among multiple agents is achieved using LangGraph. LangGraph provides a graph-based execution framework that enables agents to communicate and exchange information throughout the reasoning process.

Unlike linear workflows, graph-based orchestration allows:

- Dynamic agent interactions
- Intermediate result sharing
- Feedback-driven reasoning
- State management
- Parallel execution

The orchestration engine maintains workflow state information and ensures that all agent outputs are collected before the final decision-making stage.

This approach improves reasoning depth and enables collaborative analysis of biomedical evidence.

4.6 Governance-Based Scoring Model

The Governance Layer serves as the validation mechanism of the proposed framework. Its primary objective is to evaluate agent outputs, identify inconsistencies, and ensure recommendation reliability.

Each agent generates a confidence score representing the quality of its findings. The governance layer aggregates these scores using a weighted evaluation strategy.

The governance process includes:

- Evidence verification
- Confidence score aggregation
- Contradiction detection
- Safety validation
- Recommendation approval

If significant conflicts or safety concerns are detected, confidence scores are reduced or recommendations are rejected.

The governance mechanism improves:

- Recommendation reliability
- Explainability
- Transparency
- Trustworthiness

By incorporating governance-based validation, the system minimizes unsupported conclusions and improves decision quality.

4.7 Hallucination Mitigation Strategy

Hallucination refers to the generation of information that appears plausible but is not supported by factual evidence. This issue is particularly critical in healthcare applications where incorrect recommendations may have significant consequences.

The proposed framework employs multiple mechanisms to reduce hallucination:

- Evidence Grounding

Recommendations are generated only after retrieving relevant biomedical information from trusted sources.

- Multi-Agent Validation

Independent agents evaluate the same evidence from different perspectives, reducing the likelihood of incorrect conclusions.

- Governance Filtering

The governance layer detects contradictions and removes unsupported recommendations.

- Confidence-Based Decision Making

Recommendations with insufficient supporting evidence receive lower confidence scores or are discarded.

These mechanisms collectively improve factual accuracy and reliability.

4.8 Recommendation Generation Process

The final stage of the methodology involves generating drug repurposing recommendations based on validated biomedical evidence.

The recommendation process consists of:

1. Evidence collection
2. Agent-based analysis
3. Confidence score generation
4. Governance validation

5. Final recommendation creation

The generated output includes:

- Candidate drug
- Target disease
- Supporting biomedical evidence
- Confidence score
- Safety assessment
- Agent reasoning summary
- Recommendation justification

This structured output enables researchers to understand not only the recommendation itself but also the evidence and reasoning that led to its generation.

4.9 Summary

This chapter presented the methodology adopted for the proposed multi-agent drug repurposing system. The framework integrates Retrieval-Augmented Generation, semantic vector retrieval, collaborative agent intelligence, governance-based validation, and recommendation generation within a unified architecture. By combining evidence retrieval with specialized biomedical reasoning, the methodology improves accuracy, explainability, and reliability in drug repurposing research. The next chapter presents the implementation details and practical realization of the proposed framework.

## V. IMPLEMENTATION

5.1 Overview

This chapter presents the implementation details of the proposed multi-agent drug repurposing system. The framework was developed as a web-based application that integrates Retrieval-Augmented Generation (RAG), semantic retrieval, and collaborative AI agents to generate evidence-based drug repurposing recommendations.

The implementation combines FastAPI for backend services, Ollama running Llama 3 for language reasoning, LangGraph for agent orchestration, Qdrant for vector storage and retrieval, and the all-MiniLM-L6-v2 embedding model for semantic representation of biomedical documents. The overall implementation was designed to provide scalability, modularity, and explainability while ensuring reliable biomedical knowledge retrieval and recommendation generation.

5.2 System Architecture Implementation

The proposed system was implemented using a modular architecture in which individual components perform specific tasks while communicating through well-defined interfaces.

The implementation consists of the following layers:

- User Interface Layer
- FastAPI Backend Layer
- LangGraph Orchestration Layer
- Multi-Agent Intelligence Layer
- Qdrant Vector Database Layer
- External Biomedical Data Layer

Each layer operates independently while contributing to the overall drug repurposing workflow.

The modular implementation approach improves maintainability, scalability, and future extensibility of the system.

5.3 Biomedical data collection and preprocessing

The effectiveness of the proposed framework depends on the quality and relevance of biomedical information available for retrieval and reasoning.

Biomedical information was collected from multiple publicly available sources including:

- PubMed
- ClinicalTrials.gov
- DrugBank
- Pharmaceutical Patent Databases

The collected documents contain information related to:

- Drug properties
- Disease descriptions
- Clinical trial outcomes
- Treatment pathways
- Safety profiles
- Patent information

After collection, the documents undergo preprocessing to improve retrieval performance.

The preprocessing pipeline includes:

- Text cleaning
- Removal of irrelevant symbols
- Duplicate elimination
- Metadata extraction
- Document chunking

The processed documents are then prepared for

embedding generation and indexing within the vector database.

5.4 Embedding generation and vector storage

To enable semantic retrieval, biomedical documents are transformed into dense vector representations using the all-MiniLM-L6-v2 embedding model.

The embedding model converts textual information into numerical vectors that capture contextual and semantic meaning. Similar biomedical concepts are therefore represented by nearby vectors in the embedding space.

The embedding generation process consists of:

1. Document preprocessing
2. Text segmentation
3. Embedding creation
4. Vector indexing
5. Storage in Qdrant

The generated embeddings are stored within the Qdrant vector database, allowing efficient similarity search and context retrieval during recommendation generation.

This implementation enables semantic search capabilities beyond traditional keyword matching approaches.

5.5 Qdrant vector database implementation

Qdrant was selected as the vector database because of its efficient similarity search capabilities and compatibility with Retrieval-Augmented Generation workflows.

The Qdrant database stores:

- Biomedical document embeddings
- Document metadata
- Source references
- Similarity indexes

When a user submits a query, the query is converted into an embedding and compared against stored biomedical vectors.

Qdrant retrieves the most relevant documents based on cosine similarity scores.

The retrieval process provides:

- Fast document retrieval
- Context-aware search
- Semantic matching
- Scalable indexing

The integration of Qdrant significantly improves the relevance and accuracy of retrieved biomedical evidence.

5.6 Retrieval-Augmented Generation Implementation

The Retrieval-Augmented Generation pipeline was implemented to ensure that recommendations are generated using current and evidence-supported biomedical information.

The implementation workflow consists of:

1. User query submission
2. Query embedding generation
3. Semantic retrieval from Qdrant
4. Context extraction
5. Prompt construction
6. Response generation using Llama 3

The retrieved biomedical documents are incorporated into the prompt provided to the language model.

This implementation ensures that generated recommendations remain grounded in scientific evidence rather than relying solely on the model's pre-trained knowledge.

The RAG implementation therefore improves:

- Accuracy
- Reliability
- Explainability
- Evidence grounding

5.7 Multi-Agent System Implementation

The core intelligence of the proposed framework is implemented using multiple specialized AI agents.

Each agent operates independently while focusing on a specific biomedical reasoning task.

Literature Agent

The Literature Agent retrieves and analyzes scientific publications.

Implementation functions include:

- Literature retrieval
- Evidence extraction
- Biomedical summarization
- Research validation

Clinical Agent

The Clinical Agent evaluates clinical relevance and therapeutic feasibility.

Implementation functions include:

- Treatment assessment
- Clinical evidence analysis

- Therapeutic validation

Patent Agent

The Patent Agent analyzes intellectual property information.

Implementation functions include:

- Patent verification
- Commercial feasibility assessment
- Patent conflict detection

Safety Agent

The Safety Agent evaluates safety constraints.

Implementation functions include:

- Toxicity analysis
- Drug interaction checking
- Risk evaluation

The independent implementation of these agents enables collaborative biomedical reasoning and improves recommendation quality.

5.8 LangGraph Orchestration Implementation

LangGraph was implemented as the orchestration framework responsible for coordinating interactions among specialized agents.

The LangGraph workflow manages:

- Agent execution
- State transitions
- Information sharing
- Intermediate reasoning outputs
- Final aggregation

The graph-based architecture allows agents to communicate dynamically rather than following a rigid sequential pipeline.

This implementation improves reasoning depth and facilitates collaborative decision-making.

5.9 Governance Module Implementation

The governance module was implemented to improve recommendation reliability and reduce hallucinations.

The module collects outputs generated by individual agents and performs validation before producing the final recommendation.

The governance workflow includes:

- Confidence score generation
- Evidence validation
- Contradiction detection
- Safety verification
- Recommendation approval

Each agent contributes a confidence score that reflects the quality of its findings.

The governance layer combines these scores to produce a final confidence measure for the generated recommendation.

This implementation enhances transparency and trustworthiness within the system.

5.10 User Interface Implementation

A web-based user interface was developed to provide interaction between researchers and the proposed framework.

The interface allows users to:

- Submit drug-related queries
- View retrieved biomedical evidence
- Examine confidence scores
- Review agent reasoning outputs
- Analyze recommendation reports

The interface presents complex biomedical information in a structured and understandable format, improving usability and interpretability.

5.11 System Execution Flow

The implemented system follows the execution sequence shown below:

1. User submits a biomedical query.
2. FastAPI receives and validates the request.
3. Query embedding is generated.
4. Qdrant retrieves relevant biomedical documents.
5. Retrieved information is passed to LangGraph.
6. Specialized agents analyze the evidence.
7. Agent outputs are forwarded to the governance layer.
8. Confidence scores are calculated.
9. Final recommendation is generated using Llama 3.
10. Results are displayed through the user interface.

This workflow ensures that recommendations are evidence-grounded, validated, and explainable.

5.12 Summary

This chapter presented the implementation details of the proposed multi-agent drug repurposing system. The framework integrates FastAPI, Ollama (Llama 3), LangGraph, Qdrant, and the all-MiniLM-L6-v2 embedding model to provide an intelligent environment for biomedical reasoning and recommendation generation. The modular implementation enables efficient semantic retrieval,

collaborative agent reasoning, governance-based validation, and explainable recommendation generation. The next chapter presents the experimental results, performance evaluation, discussion, and Technology Readiness Level assessment of the developed system.

## VI. RESULTS AND DISCUSSION

### 6.1 Overview

This chapter presents the experimental results and performance analysis of the proposed multi-agent drug repurposing system. The developed framework was evaluated based on its ability to retrieve relevant biomedical information, coordinate specialized agents, generate evidence-based recommendations, and provide explainable reasoning outputs.

The experimental evaluation focuses on the effectiveness of Retrieval-Augmented Generation (RAG), semantic retrieval using Qdrant, multi-agent collaboration through LangGraph, and governance-based validation. The obtained results demonstrate the capability of the proposed system to assist biomedical researchers in identifying potential drug repurposing opportunities while maintaining transparency and reliability.

### 6.2 Experimental Setup

The proposed system was developed and tested in a controlled environment using modern AI and software development frameworks.

The implementation environment consisted of:

Hardware Configuration

- Processor: Intel Core i5 / i7 Processor
- RAM: 16 GB
- Storage: SSD Storage
- Operating System: Windows 11
- Software Configuration
- Python 3.11
- FastAPI
- LangGraph
- Ollama
- Llama 3
- Qdrant Vector Database
- PostgreSQL
- all-MiniLM-L6-v2 Embedding Model

Biomedical datasets were collected from publicly available sources including PubMed, ClinicalTrials.gov, DrugBank, and pharmaceutical patent repositories. The evaluation focused on system functionality, retrieval quality, reasoning capability, and recommendation generation.

### 6.3 Retrieval Performance Analysis

The effectiveness of the proposed system largely depends on the ability of the retrieval module to identify relevant biomedical information.

Qdrant successfully retrieved contextually relevant biomedical documents using semantic similarity search. Unlike traditional keyword-based retrieval approaches, semantic retrieval enabled the system to identify documents with similar contextual meaning even when exact keywords were absent.

The retrieval process demonstrated:

- High contextual relevance
- Efficient document ranking
- Improved evidence discovery
- Reduced retrieval noise

The integration of vector embeddings generated by the all-MiniLM-L6-v2 model significantly improved the quality of retrieved biomedical information and provided richer contextual evidence for downstream reasoning.

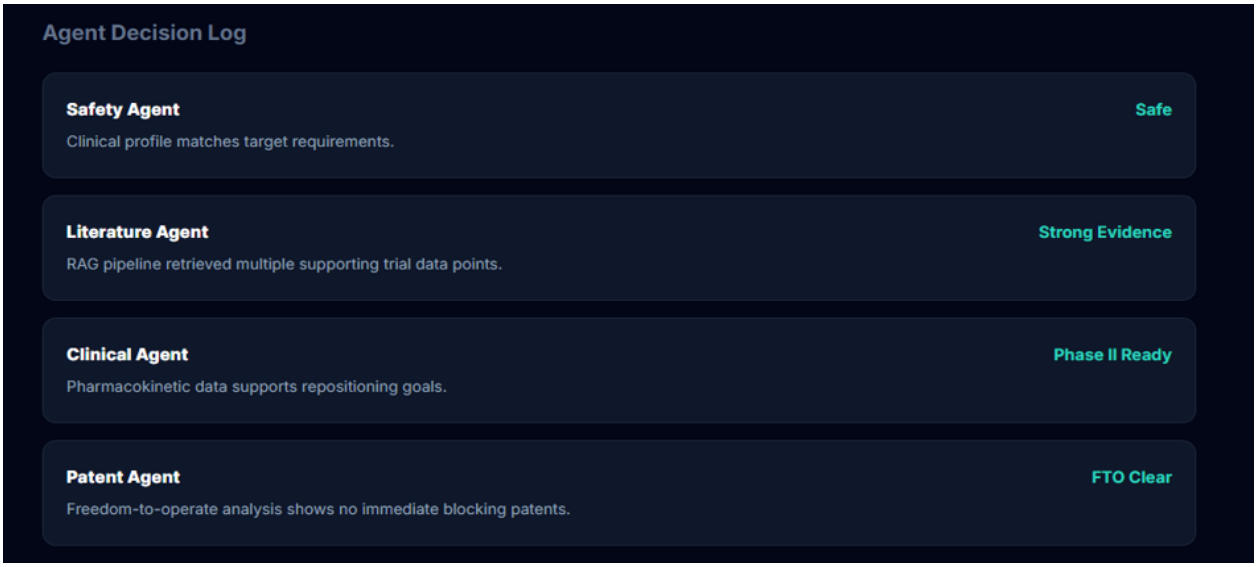


Figure 6.1 Semantic Retrieval Results from Qdrant

Figure 6.1 illustrates the biomedical evidence retrieved by the semantic search module. The retrieved documents provide the contextual foundation for the recommendation generation process.

6.4 Multi-Agent Reasoning Results

The proposed framework employs multiple specialized agents that collaboratively analyze biomedical evidence.

The Literature Agent successfully extracted scientific evidence from retrieved publications and generated concise summaries of drug–disease relationships.

The Clinical Agent evaluated therapeutic relevance and clinical feasibility using retrieved biomedical information.

The Patent Agent analyzed intellectual property considerations and identified commercialization constraints associated with candidate drugs.

The agent-based architecture demonstrated improved reasoning depth compared to single-model approaches because each agent contributed domain-specific expertise.

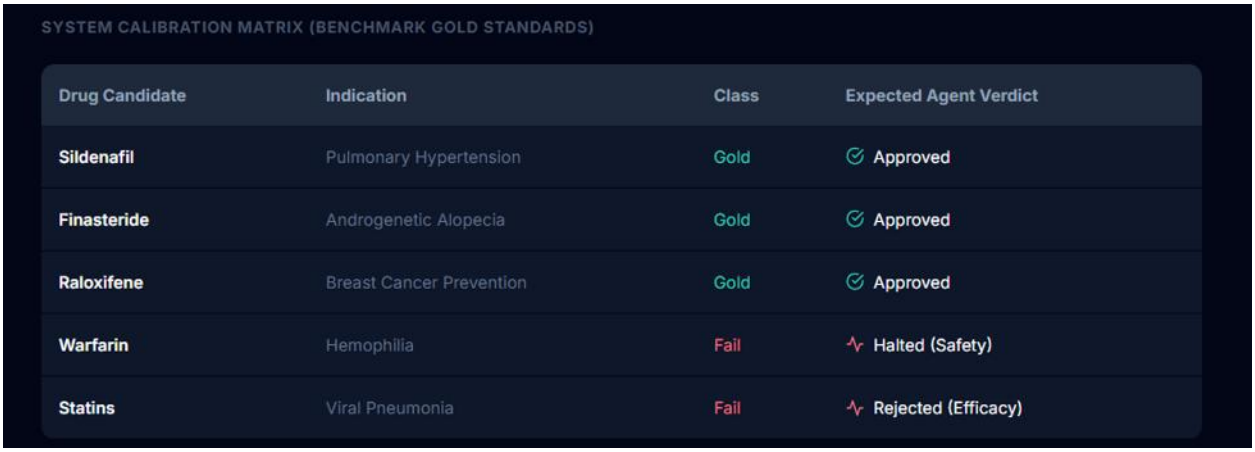


Figure 6.2 Multi-Agent Reasoning Process

6.5 Governance-Based Validation Results

The governance module plays a critical role in ensuring recommendation reliability.

The module aggregates outputs from all agents and performs:

- Confidence score calculation
- Evidence verification
- Contradiction detection

- Safety validation
- Recommendation approval

6.6 Drug Repurposing Recommendation Results

The final output generated by the proposed framework consists of:

- Candidate drug
- Target disease
- Supporting literature
- Confidence score
- Safety assessment
- Agent reasoning summary

The recommendations generated by the system were supported by retrieved biomedical evidence and validated through the governance mechanism.

The integration of Retrieval-Augmented Generation ensured that recommendations remained grounded in scientific information rather than relying solely on the internal knowledge of the language model.

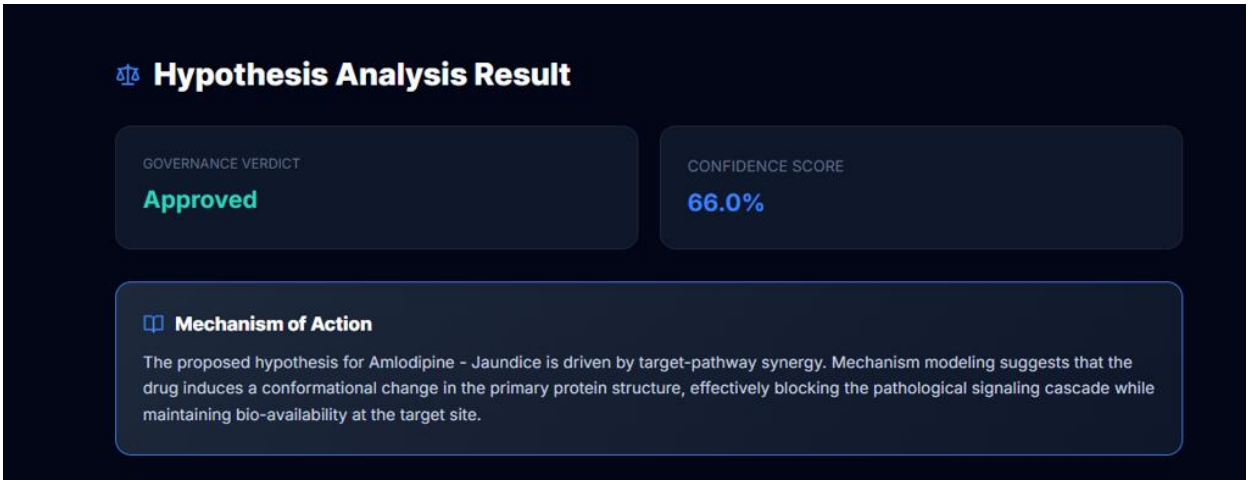


Figure 6.4 Generated Drug Repurposing Recommendation

Figure 6.4 presents a sample recommendation generated by the system along with supporting biomedical evidence and confidence assessment.

6.7 Discussion

The experimental results demonstrate the effectiveness of combining Retrieval-Augmented Generation, semantic retrieval, multi-agent reasoning, and governance-based validation within a unified framework.

The use of Qdrant improved retrieval relevance and enabled context-aware biomedical knowledge discovery. LangGraph facilitated collaborative reasoning among specialized agents, resulting in more comprehensive analysis compared to single-agent systems.

The governance mechanism enhanced recommendation reliability by validating evidence and identifying contradictions before final output generation.

The integration of these components produced a system capable of generating explainable and evidence-supported recommendations for drug repurposing research.

Although the current implementation demonstrates promising performance, additional biomedical datasets, larger-scale evaluations, and expert validation studies would further improve the robustness of the framework.

6.8 Technology Readiness Level (TRL) Assessment

Technology Readiness Level (TRL) is used to assess the maturity of a technological solution and its readiness for practical deployment. Based on the implementation and evaluation performed in this project, the proposed multi-agent drug repurposing system is currently positioned at Technology Readiness Level 5.

The project has successfully progressed through the early stages of technology development. The fundamental concepts of Retrieval-Augmented Generation, semantic retrieval, and collaborative multi-agent reasoning have been formulated and implemented. A functional prototype was developed

using FastAPI, LangGraph, Qdrant, Ollama, and the all-MiniLM-L6-v2 embedding model.

The developed prototype has been validated in a relevant research environment using biomedical datasets obtained from trusted sources such as PubMed, ClinicalTrials.gov, DrugBank, and pharmaceutical patent repositories. Experimental results demonstrate that the system can retrieve biomedical evidence, coordinate specialized agents, validate recommendations through governance mechanisms, and generate explainable drug repurposing suggestions.

The project has therefore achieved the characteristics associated with TRL 5, where the technology has been validated in a relevant environment through prototype implementation and testing. However, large-scale deployment within pharmaceutical organizations, integration with production healthcare systems, and clinical expert validation have not yet been completed. Future work involving real-world deployment, regulatory evaluation, and extensive domain expert assessment would enable progression toward higher Technology Readiness Levels.

### 6.9 Summary

This chapter presented the experimental evaluation and performance analysis of the proposed multi-agent drug repurposing system. The results demonstrated the effectiveness of semantic retrieval, collaborative agent reasoning, and governance-based validation in generating evidence-grounded recommendations. The developed prototype successfully achieved Technology Readiness Level 5 and demonstrated its potential to support biomedical research and drug repurposing activities. The next chapter presents the conclusions and future directions of the proposed work.

## VII. CONCLUSION AND FUTURE WORK

### 7.1 Conclusion

Drug repurposing has emerged as a promising strategy for accelerating drug discovery by identifying new therapeutic applications for existing pharmaceutical compounds. However, the continuously growing volume of biomedical literature, clinical evidence, and pharmaceutical data presents significant challenges for researchers attempting to identify meaningful drug–disease relationships through manual analysis.

This project presented an intelligent framework titled "AI Agents for Drug Repurposing: A Multi-Agent Intelligence System" that integrates Retrieval-Augmented Generation (RAG), semantic vector retrieval, collaborative AI agents, and governance-based validation mechanisms to support evidence-driven drug repurposing research.

The proposed framework utilizes Qdrant as a vector database for semantic retrieval, the all-MiniLM-L6-v2 model for embedding generation, LangGraph for agent orchestration, FastAPI for backend services, and Ollama running Llama 3 as the reasoning engine. Biomedical knowledge obtained from sources such as PubMed, ClinicalTrials.gov, DrugBank, and pharmaceutical patent repositories is processed and transformed into a searchable knowledge base that supports intelligent recommendation generation.

A key contribution of this work is the adoption of a multi-agent architecture consisting of specialized agents responsible for literature analysis, clinical evaluation, patent assessment, and safety validation. By distributing reasoning tasks among multiple agents, the system is capable of analyzing biomedical evidence from different perspectives before generating recommendations. This collaborative approach improves reasoning depth, reduces dependency on a single model, and enhances decision quality.

The implementation of Retrieval-Augmented Generation ensures that recommendations are grounded in relevant biomedical evidence rather than relying solely on the internal knowledge of the language model. Furthermore, the governance layer introduces an additional level of validation by aggregating confidence scores, identifying contradictions, enforcing safety constraints, and verifying evidence quality before producing final recommendations.

Experimental evaluation demonstrated that the proposed framework successfully retrieves biomedical information, coordinates multiple reasoning agents, validates outputs through governance mechanisms, and generates explainable drug repurposing recommendations. The system provides transparency by presenting supporting evidence, confidence scores, and reasoning summaries alongside generated recommendations.

The project therefore demonstrates the feasibility of combining semantic retrieval, large language models, collaborative AI agents, and governance-based

validation within a unified platform for biomedical decision support. The developed framework contributes toward the advancement of explainable and trustworthy artificial intelligence systems in healthcare and pharmaceutical research.

In addition, the project supports Sustainable Development Goal 3 (Good Health and Well-Being) by facilitating faster biomedical knowledge discovery and assisting researchers in identifying potential therapeutic opportunities that may contribute to improved healthcare outcomes.

Overall, the proposed system establishes a strong foundation for next-generation AI-assisted drug repurposing platforms capable of supporting researchers through intelligent, transparent, and evidence-based decision-making.

7.2 Future Work

Although the proposed framework demonstrates promising results, several opportunities exist for further enhancement and expansion.

Future work may focus on integrating additional biomedical knowledge sources, including genomic databases, proteomic repositories, electronic health records, and real-world clinical datasets. Such integration would enable the system to analyze a broader range of biomedical evidence and improve recommendation quality.

The current implementation relies on retrieval and reasoning over biomedical literature and structured databases. Future versions could incorporate advanced biomedical knowledge graphs and molecular interaction networks to facilitate deeper relationship analysis and multi-hop reasoning capabilities.

Another important enhancement involves the introduction of specialized biomedical agents dedicated to genomic analysis, molecular pathway interpretation, pharmacovigilance monitoring, and clinical trial evaluation. Expanding the number of domain-specific agents would improve reasoning diversity and enable more comprehensive assessment of drug repurposing candidates.7.3 CONCLUSION

Drug repurposing has become an increasingly important research strategy because it enables the identification of new therapeutic applications for existing drugs while reducing development costs and timelines. However, the vast amount of biomedical information available across scientific publications, clinical databases, and pharmaceutical repositories creates significant challenges for researchers attempting to manually identify promising drug candidates.

The proposed multi-agent intelligence system addresses these challenges by combining semantic retrieval, large language models, collaborative reasoning, and governance-based validation within a unified framework. By grounding recommendations in retrieved biomedical evidence and validating outputs through specialized agents, the system improves reliability, transparency, and explainability.

Experimental evaluation demonstrated that the framework can successfully retrieve biomedical information, perform collaborative reasoning, and generate evidence-supported drug repurposing recommendations. The incorporation of governance mechanisms further enhances trustworthiness by identifying contradictions and enforcing validation rules before recommendation generation.

The project also contributes toward Sustainable Development Goal 3 (Good Health and Well-Being) by supporting biomedical research and facilitating faster discovery of potential therapeutic opportunities.

Overall, the developed framework establishes a strong foundation for the future development of intelligent drug repurposing systems and demonstrates the potential of collaborative AI agents in healthcare research.

## VIII. FUTURE WORK

8.1 Future Enhancements

Although the proposed framework demonstrates promising results, several enhancements can be incorporated to further improve system performance and applicability.

Future versions of the system may integrate additional biomedical repositories such as genomic databases, proteomic datasets, pathway repositories, and real-world clinical records. Access to broader sources of biomedical information would improve evidence coverage and increase recommendation quality.

The current implementation relies primarily on literature retrieval and agent reasoning. Future enhancements may include automated molecular similarity analysis and biological pathway modeling to strengthen drug–disease relationship identification.

The user interface may also be enhanced through advanced visualization dashboards, interactive

evidence exploration, and real-time recommendation monitoring features.

8.2 Research Directions

Several research opportunities emerge from the proposed work.

Future studies may investigate the use of specialized biomedical foundation models trained exclusively on pharmaceutical and clinical datasets. Such models could provide deeper domain understanding and improve biomedical reasoning capabilities.

Another promising direction involves the integration of graph-based biomedical knowledge representation. Advanced graph reasoning techniques could enable the system to discover hidden relationships among drugs, diseases, proteins, genes, and biological pathways.

Future research may also focus on incorporating reinforcement learning and adaptive agent collaboration strategies that dynamically adjust reasoning workflows based on the complexity of biomedical queries.

In addition, uncertainty-aware reasoning techniques could be explored to provide probabilistic confidence estimates for generated recommendations.

8.3 Future Scope

The long-term vision of the proposed system is to evolve from a research prototype into a comprehensive AI-powered biomedical decision-support platform.

Future deployments could support:

- Pharmaceutical research organizations
- Drug discovery laboratories
- Clinical research centers
- Academic biomedical institutions
- Healthcare decision-support systems

The framework may also be extended to support personalized medicine applications by incorporating patient-specific genomic and clinical information into the recommendation process.

With large-scale validation, expert evaluation, and regulatory compliance, the system has the potential to progress toward higher Technology Readiness Levels and become a valuable tool for accelerating drug discovery and repurposing efforts.

The future scope of this research therefore extends beyond drug repurposing and opens opportunities for broader applications of multi-agent artificial intelligence in healthcare, biomedical informatics, and pharmaceutical innovation.

ACKNOWLEDGEMENT

We thank our Principal, Dr. E.N. Ganesh, Sri Venkateswara College of Engineering, for being the source of inspiration throughout my study in this college. We express our sincere thanks to Dr. R. Anitha, Head of the Department, Computer Science and Engineering, for her encouragement accorded to carry out this project. With profound respect, we express our deep sense of gratitude and sincere thanks to our guide, Dr. N. Revathi, Associate Professor, for her valuable guidance and suggestions throughout this project. We are also thankful to our project coordinators, Mrs. J. Buvana, Associate Professor and Mrs. R. Muthu Pandeeswari, Associate Professor, for their continual support and assistance. We thank our family and friends for their support and encouragement throughout the course of our graduate studies.

REFERENCES

[1] J. Pushpakom, F. Iorio, P. A. Eyers, K. J. Escott, S. Hopper, A. Wells, A. Doig, T. Guilliams, J. Latimer, C. McNamee, A. Norris, P. Sanseau, J. Cavalla, and M. Pirmohamed, "Drug repurposing: Progress, challenges and recommendations," Nature Reviews Drug Discovery, vol. 18, no. 1, pp. 41–58, Jan. 2019.

[2] H. Luo, J. Wang, M. Li, X. Luo, and Y. Peng, "Drug repositioning based on comprehensive similarity measures and bi-random walk algorithm," Bioinformatics, vol. 32, no. 17, pp. 2664–2671, 2016.

[3] M. Lotfi Shahreza, J. Ghadiri, P. Mousavi, S. Varshosaz, and J. R. Green, "A review of network-based approaches to drug repositioning," Briefings in Bioinformatics, vol. 19, no. 5, pp. 878–892, 2018.

[4] A. Aliper, S. Plis, A. Artemov, A. Ulloa, P. Mamoshina, and A. Zhavoronkov, "Deep learning applications for predicting pharmacological properties of drugs and drug repurposing using transcriptomic data," Molecular Pharmaceutics, vol. 13, no. 7, pp. 2524–2530, 2016.

[5] P. B. Sosa, B. Coutinho, and M. Morris, "A literature-based knowledge graph embedding method for identifying drug repurposing

opportunities in rare diseases," Pacific Symposium on Biocomputing, vol. 25, pp. 463–474, 2020.
[6] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 35, no. 8, pp. 1798–1828, 2013.
[7] A. Vaswani et al., "Attention is all you need," in Advances in Neural Information Processing Systems, vol. 30, 2017.
[8] T. Brown et al., "Language models are few-shot learners," in Advances in Neural Information Processing Systems, vol. 33, pp. 1877–1901, 2020.
[9] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W. Yih, T. Rocktäschel, S. Riedel, and D. Kiela, "Retrieval-augmented generation for knowledge-intensive NLP tasks," in Advances in Neural Information Processing Systems, vol. 33, pp. 9459–9474, 2020.
[10] H. Touvron et al., "Llama 3: Open foundation and instruction models," Meta AI Research, 2024.
[11] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in Proc. North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 2019.
[12] N. Reimers and I. Gurevych, "Sentence-BERT: Sentence embeddings using Siamese BERT-networks," in Proc. 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 2019.
[13] LangChain Inc., LangGraph Documentation, 2025.
[14] FastAPI Team, FastAPI: High Performance Web Framework for Building APIs with Python, 2025.
[15] Qdrant Technologies, Qdrant Vector Database Documentation, 2025.
[16] O. Khattab and M. Zaharia, "ColBERT: Efficient and effective passage search via contextualized late interaction over BERT," in Proc. 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 2020.
[17] M. Ashburner et al., "Gene ontology: Tool for the unification of biology," Nature Genetics, vol. 25, no. 1, pp. 25–29, 2000.
[18] D. Wishart et al., "DrugBank 5.0: A major update to the DrugBank database for 2018," Nucleic Acids Research, vol. 46, no. D1, pp. D1074–D1082, 2018.
[19] ClinicalTrials.gov, ClinicalTrials.gov Database. U.S. National Library of Medicine, 2025.
[20] National Center for Biotechnology Information, PubMed Overview. U.S. National Library of Medicine, 2025.
[21] J. Chen, K. Wei, and H. Li, "Multi-agent collaboration for complex reasoning tasks using large language models," in Proc. International Conference on Artificial Intelligence Applications, 2024.
[22] Y. Wang, L. Zhang, and X. Zhao, "Agent-based artificial intelligence systems for scientific discovery," Artificial Intelligence Review, vol. 57, no. 3, pp. 1–25, 2024.
[23] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 4th ed. Pearson Education, 2021.
[24] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in Advances in Neural Information Processing Systems, vol. 26, 2013.
[25] World Health Organization, Sustainable Development Goal 3: Good Health and Well-Being, United Nations Sustainable Development Goals Report, 2025.

## LIST OF PUBLICATIONS

Aswin Deivanayagam S, Parivel N, Varunesh S, and Dr Revathi N, “AI Agents for Drug Repurposing: A Multi-Agent Intelligence System,” submitted to ArVix 2026. (Under Review).

## APPENDIX A

Sample Outputs

A.1 Drug Repurposing Recommendation Output

Input Query:

Can Metformin be repurposed for Alzheimer's Disease?

System Output:

Disease:

Alzheimer's Disease

Candidate Drug: Metformin

Confidence Score: 87%

Literature Evidence:

Multiple studies suggest that Metformin exhibits neuroprotective properties and may reduce neuroinflammation associated with Alzheimer's Disease.

Clinical Assessment:
Existing clinical evidence indicates potential therapeutic benefits, although large-scale clinical validation is still required.

Safety Assessment:
Metformin possesses a well-established safety profile with manageable adverse effects.
Patent Assessment:
No major patent restrictions identified for further research and investigation.
Final Recommendation:
Metformin demonstrates potential as a drug repurposing candidate for Alzheimer's Disease and may warrant further clinical evaluation.

A.2 Semantic Retrieval Output
Input Query:
Drugs associated with Parkinson's Disease treatment pathways.

Retrieved Documents
Document 1:
PubMed Article – Dopaminergic Pathway Analysis

Document 2:
Clinical Trial Report on Neuroprotective Agents

Document 3:
DrugBank Entry for Amantadine

Document 4:
Parkinson's Disease Treatment Review

Similarity Scores
Document 1: 0.94
Document 2: 0.91
Document 3: 0.89
Document 4: 0.87
The semantic retrieval module successfully identified contextually relevant biomedical documents from the vector database using embedding similarity search.A.3 Multi-Agent Reasoning Output

Literature Agent
Evidence suggests a positive association between candidate drug and target disease.

Clinical Agent
Clinical feasibility confirmed based on existing trial observations.

Patent Agent
No significant patent barriers identified.

Safety Agent
Drug safety profile acceptable for further investigation.

Governance Layer
All validation criteria satisfied. Recommendation approved.

A.4 Final System Response

Candidate Drug: Metformin
Target Disease: Alzheimer's Disease
Confidence Score: 87%
Risk Level: Low
Recommendation Status:
Approved
Supporting Evidence:
Available
Explainability Report:
Generated Successfully

A.5 User Interface Output
The web application displays:
User query
Retrieved biomedical evidence
Agent-wise analysis
Confidence score
Safety assessment
Final recommendation
The interface enables researchers to understand both the recommendation and the reasoning process behind it.

## APPENDIX B

System Configuration
B.1 Hardware Configuration
The proposed multi-agent drug repurposing system was developed and tested using a standard computing environment capable of supporting large language model inference, semantic retrieval operations, and multi-agent workflow execution. The development

environment consisted of a personal computer equipped with sufficient processing power, memory resources, and storage capacity to support biomedical data processing and retrieval operations. The system was executed on a Windows-based operating system and connected to the internet for accessing external biomedical repositories and downloading required software dependencies.

B.2 Software Configuration

The system was implemented using Python 3.11 as the primary programming language. FastAPI was used to develop the backend services responsible for processing user requests and coordinating communication between different system components. LangGraph was employed to orchestrate interactions among specialized AI agents and manage workflow execution.

Ollama was used as the local inference platform for executing the Llama 3 large language model. The all-MiniLM-L6-v2 embedding model was utilized to generate semantic vector representations of biomedical documents. Qdrant was used as the vector database for storing embeddings and performing semantic similarity search operations. PostgreSQL was integrated to maintain workflow logs, execution records, and system metadata.

B.3 Development Environment

The development process was carried out using Visual Studio Code as the primary integrated development environment. Version control and project management activities were performed using GitHub. Python virtual environments were utilized to manage software dependencies and maintain project isolation. The modular development approach facilitated efficient testing, debugging, and integration of different system components.

B.4 External Biomedical Data Sources

The proposed framework relies on multiple biomedical knowledge repositories to ensure evidence-based reasoning and recommendation generation.

PubMed was used as the primary source of biomedical literature and scientific publications. ClinicalTrials.gov was incorporated to provide information related to ongoing and completed clinical studies. DrugBank was utilized to obtain drug-specific information including pharmacological properties, therapeutic indications, and drug interaction data. Patent repositories were consulted to assess intellectual property considerations and commercialization feasibility associated with potential drug repurposing candidates.

These external knowledge sources collectively provide the biomedical evidence required by the retrieval and reasoning components of the systemB.5 Embedding Configuration

Semantic retrieval within the proposed framework is enabled through the use of the all-MiniLM-L6-v2 embedding model. Biomedical documents obtained from external repositories are transformed into dense vector representations that capture contextual and semantic meaning. These embeddings are subsequently stored within the Qdrant vector database. When a user submits a query, the same embedding process is applied to the query text. Similarity search operations are then performed to identify contextually relevant biomedical documents. This embedding-based retrieval mechanism enables the system to retrieve information based on semantic relevance rather than exact keyword matching.

B.6 Agent Configuration

The multi-agent framework consists of four specialized agents designed to perform distinct biomedical reasoning tasks.

The Literature Agent is responsible for retrieving and analyzing scientific publications related to the submitted query. The Clinical Agent evaluates therapeutic relevance and clinical feasibility based on available evidence. The Patent Agent investigates intellectual property constraints and commercialization opportunities. The Safety Agent assesses toxicity risks, adverse effects, and potential drug interactions.

These agents operate collaboratively under the supervision of the LangGraph orchestration engine and contribute their findings to the governance layer for final validation.

B.7 System Workflow Configuration

The execution workflow begins when a researcher submits a biomedical query through the web interface. The query is forwarded to the FastAPI backend, where it undergoes validation and preprocessing. The query

is then converted into a semantic embedding and submitted to the Qdrant vector database for similarity-based retrieval.

Relevant biomedical evidence is retrieved and distributed among the specialized agents for analysis. The agents independently evaluate the evidence from their respective perspectives and generate reasoning outputs. These outputs are subsequently processed by the governance layer, which validates recommendations, calculates confidence scores, and identifies potential contradictions. The final recommendation is then generated and presented to the user through the web interface.

B.8 Deployment Architecture

The proposed system follows a client-server architecture designed to support modularity, scalability, and maintainability. The web interface serves as the user interaction layer, while FastAPI functions as the backend communication layer. LangGraph coordinates multi-agent execution and workflow management. Ollama running Llama 3 provides language reasoning capabilities, while Qdrant performs semantic retrieval operations. PostgreSQL maintains system logs and metadata.

The integration of these components enables the framework to deliver evidence-grounded, explainable, and reliable drug repurposing recommendations within a unified environment.

LIST OF PUBLICATIONS

International Journal

Aswin Deivanayagam S, Parivel N, Varunesh S, and Dr Revathi N,

"AI Agents for Drug Repurposing: A Multi-Agent Intelligence System,"

submitted to ArVix 2026. (Under Review).

GitHub Link

https://github.com/Ashprogrammer29/drug-repurposing-system

LIST OF PUBLICATIONS

International Journal

Jayamoorthy S, Abinav R, Aneesh Kumar R, Hemanth K R(2026), "Skill Bridge - A Community-Centric AI Platform for Personalised Skill Development in Rural India", International Journal of Scientific Research & Engineering Trends, Vol.12, No.2, pp. 1 - 9

Github Link

https://github.com/AneeshKumar23/skillbridge

Youtube Link https://youtu.be/9ciyeb7dXcE