



Leveraging Multi-Agent AI with Vector Databases

Redefine NLP Data Analysis-Marine Piracy

¹Anvesh Reddy Minukuri,

¹VP Sr Lead Machine Learning Engineer

Abstract : A few years ago, traditional NLP methods posed challenges for data scientists, but today, AI agents are revolutionizing information analysis and retrieval, enhancing speed, intelligence, and efficiency. Recent advancements in Artificial Intelligence (AI) have significantly enhanced the capabilities of data scientists, transforming traditionally labor-intensive tasks in natural language processing (NLP) into more efficient and intelligent workflows. This study explores the integration of multi-agent AI systems and vector databases to enhance the analysis and retrieval of marine piracy data, specifically focusing on the transformation of a historical PDF-based analysis method. By combining the Phidata framework with state-of-the-art AI agents, this initiative streamlines the extraction and analysis of insights from maritime piracy reports. The goal is to demonstrate the superior capabilities of modern AI technologies in automating workflows, enhancing predictive modeling, and improving the efficiency of piracy analysis.

IndexTerms – LLMs, Agentic AI, Vector DB, Multi-Agent, RAG, PhiData, Knowledge Bases

1.INTRODUCTION

Prior to LLMs, The NLP analysis required heavy manual effort, with data scientists performing coding, preprocessing, and model building manually. The team used natural language processing (NLP) techniques, such as NLTK preprocessing, topic modeling, and clustering, to better understand the patterns in piracy attacks. Predictive models were built to forecast attack types based on various factors such as ship type, region, and time of year. While the methodology provided valuable insights, it was not fully automated and required substantial time and human involvement.

Today, with the integration of AI agents and vector databases, these steps can be automated, creating seamless workflows for handling unstructured data, clustering attacks, and building predictive models. Multi-agent AI systems significantly enhance this process by delegating specific tasks to specialized agents, streamlining workflows, and enabling real-time, adaptive learning. In this article, we will explore how the implementation of multi-agent AI applications with vector databases can revolutionize piracy analysis by enhancing query responses, predictive modeling, and overall efficiency.

Marine piracy has been a significant global issue, affecting international shipping and trade. In 2015, a team of data scientists, including the authors, conducted an in-depth study to analyze piracy incidents by leveraging structured and unstructured data sourced from the International Maritime Organization (IMO). This research, "Analyzing Marine Piracy from Structured and Unstructured Data Using SAS® Text Miner," involved manual efforts in coding, data preprocessing, and model development. Although valuable insights were gained, the process was far from fully automated and required substantial human involvement. Fast-forward to today, the advent of AI agents and vector databases has revolutionized this approach, enabling the automation of complex tasks and providing more efficient and intelligent systems for piracy analysis.

2.Background and Historical Methodology

In 2015, the analysis of marine piracy was a challenging task due to the unstructured nature of piracy reports, which were primarily in PDF format. Data scientists used natural language processing (NLP) techniques, such as topic modeling, clustering, and predictive modeling, to derive insights from these reports. However, despite the success of these methods, the process remained labor-intensive and time-consuming.

3.Current Technological Advancements: Multi-Agent AI Systems and Vector Databases

With the recent advances in multi-agent AI systems and vector databases, significant improvements have been made in automating the analysis of piracy data. In particular, multi-agent AI frameworks are revolutionizing the way data is processed and queried by delegating specific tasks to specialized agents, which work together to enhance system performance and efficiency.

Vector databases, such as PGVector, allow the storage and retrieval of semantic embeddings—vector representations of the PDF content—enabling more accurate and efficient retrieval of relevant information. The combination of multi-agent AI systems and vector databases creates a synergistic environment that significantly improves the analysis of marine piracy data. With these technologies, tasks previously requiring manual intervention, such as extracting and classifying data from piracy reports, can now be fully automated.

4. The Phidata Framework for Multi-Agent PDF Analysis

The Phidata framework serves as the backbone for this innovative multi-agent approach. It integrates various AI agents that work together to load, parse, store, and analyze PDF documents related to marine piracy. Each agent performs a specialized task, streamlining the analysis process:

- **Knowledge Base Agent:** This agent is responsible for managing the loading, parsing, and storage of PDF documents into a vector database.
- **Vector Database and Embedding Agent:** This agent handles the storage and retrieval of semantic embeddings derived from the PDF content, allowing for efficient similarity-based queries.
- **Assistant Agent:** The core conversational agent that interacts with the user, processes queries, and provides relevant responses by querying the knowledge base.
- **Storage Management Agent:** Responsible for managing session storage and keeping track of user interactions, queries, and responses.

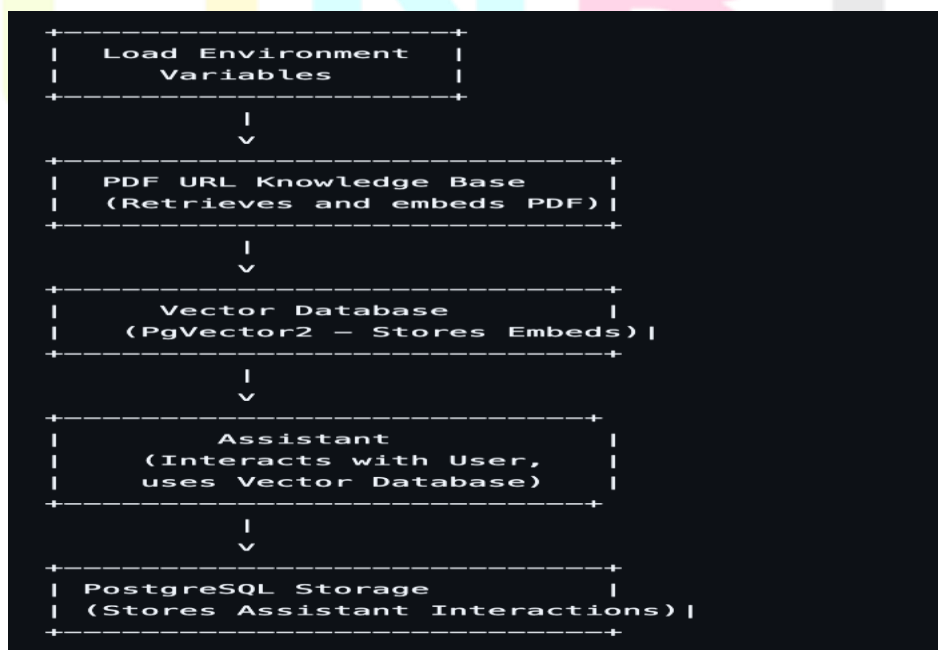
5. Implementation of Multi-Agent AI with Vector Databases

To implement the multi-agent PDF assistant for piracy analysis, the Phidata framework employs several key components (Figure below)

- **Setting Up Dependencies**
The necessary libraries, such as Typer, Phi, GroqEmbedder, and PGVector, are used to integrate AI models with the vector database and assist with embedding PDF content into vectors for later retrieval.
- **Loading Environment Variables**
Environment variables are securely loaded to ensure smooth integration with external APIs (e.g., Groq API for embedding). This ensures seamless access to embedding services and facilitates interaction with external AI models such as OpenAI.
- **Docker for PGVector Setup**
A Docker container is utilized to host PGVector, a vector database responsible for storing and managing vectorized content from the piracy reports. This allows for scalable and efficient data management.
- **Knowledge Base and Assistant Setup**
The knowledge base is populated with URLs pointing to the piracy reports, and the assistant is configured to interact with the user, process queries, and retrieve relevant information based on the loaded reports.
- **Assistant Storage and Session Management**
The assistant's state, including user queries and responses, is stored in a PostgreSQL database. This ensures continuity in user interactions, enabling the assistant to retain the context of previous sessions and adapt to evolving user queries.

The main function controls the execution of the assistant, managing user interactions and ensuring smooth operation over multiple sessions.

Figure 5.1 Multi-Agent Engineering Flow



6. Impact of Multi-Agent AI on Piracy Analysis AI

The integration of multi-agent AI and vector databases significantly improves piracy analysis by automating manual tasks that were previously required. Key improvements include:

- **Efficiency in Data Handling:** Automation of data extraction and processing enables the system to handle large volumes of piracy reports in real time, greatly reducing human intervention.
- **Improved Query Responses:** Semantic search capabilities allow users to query the system using natural language, and the assistant can generate more relevant and accurate responses.
- **Adaptive Learning:** Multi-agent systems continuously improve by learning from user interactions and new data, leading to better predictive modeling and enhanced identification of piracy patterns over time.

By reducing the time and effort involved in manual data extraction and analysis, the system provides quicker insights into piracy patterns, enabling better-informed decision-making for stakeholders in the maritime security domain.

7. Outputs: Enhanced Discovery and Insights Through Multi-Agent AI

This approach redefines the analysis of marine piracy data, leveraging multi-agent AI to automate and optimize the entire process of retrieval, analysis, insight generation, and NLP tasks. The agents not only manage data extraction and classification but also perform advanced NLP tasks and help discover hidden insights from large datasets. Below are examples of queries and outputs that illustrate the enhanced capabilities:

1. How many events are captured in the marine attacks NLP data?

- The system processes historical piracy reports and identifies the number of documented attack events stored in the vector database. The assistant queries the database for total counts, providing users with an exact figure based on semantic search of the report contents.

2. What are the different attack types captured in the knowledge base?

- The assistant extracts a list of attack types (e.g., hijacking, armed robbery, kidnapping, etc.) from the database, providing a comprehensive overview of the various types of piracy incidents stored in the knowledge base.

3. Perform topic modeling and provide the distribution of topics?

- Using NLP techniques like Latent Dirichlet Allocation (LDA), the system performs topic modeling on the available piracy reports and generates a distribution of topics (e.g., types of pirates, attack locations, motivations) along with their respective frequencies.

4. Classify attack descriptions into different categories based on attack style and pirate behavior?

- The system leverages supervised learning models and predefined categories to classify attack descriptions based on style (e.g., violent, non-violent) and pirate behavior (e.g., opportunistic, organized). The assistant provides a categorized report on the attack data.

5. Given a latitude and longitude of a ship and seashore latitudes and longitudes, and type of vessel, what is the likelihood of an attack?

- The system integrates geospatial data with historical attack patterns, assessing the likelihood of an attack based on proximity to high-risk zones and the type of vessel. The assistant computes a probability score and provides a summary of the likelihood based on the queried parameters.

These queries and outputs exemplify the power of multi-agent AI systems in providing detailed, accurate, and relevant answers to complex maritime security questions. By leveraging these technologies, the analysis of piracy incidents becomes faster, more accurate, and more adaptive to evolving threats.

8 Example Interaction Flow

Users can interact with the AI assistant by asking specific questions related to piracy incidents. The assistant processes these queries, searches the vector database for relevant information, and generates responses. For example, a user might ask about the frequency of different types of pirate attacks or inquire about common piracy behaviors in a specific region. The assistant retrieves and presents the relevant data in a user-friendly format, such as a frequency table or descriptive analysis.

The assistant processes these queries, searches the vector database for relevant information, and generates responses. For example, a user might ask (GenAI App interface outputs):

- "How many attack types are captured in the knowledge base?"
- "Give me a frequency table of attack types."
- "What are the most common pirate behaviors in the Gulf of Aden?"

Appendix:

For further implementation details, including code snippets and configuration steps, refer to the GitHub repository <https://github.com/anvcse562/Multi-Agent-AI-RAG-PDF-Marine-Privacy>

Figure 7.1 NLP numerical parsing




Figure 7.2 NLP statistical parsing

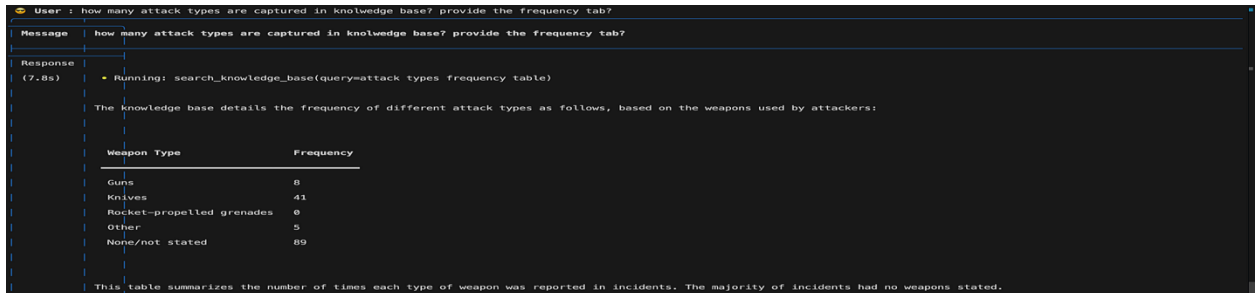
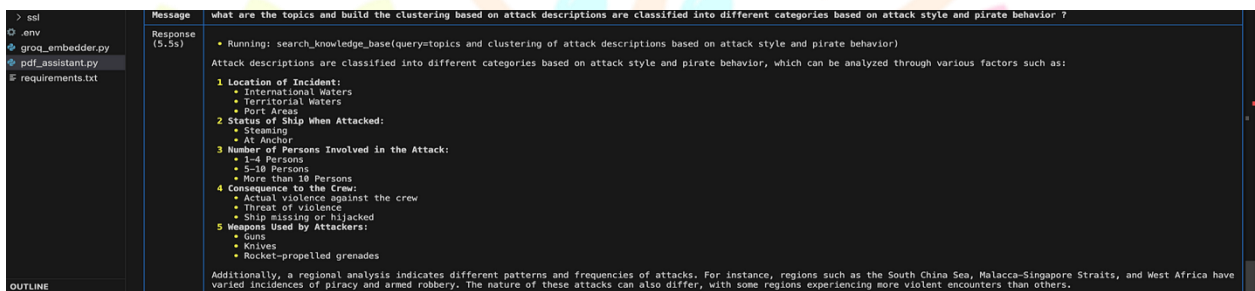


Figure 7.3 NLP tasks execution & Insights



8. Conclusion

This paper demonstrates how integrating multi-agent AI systems with vector databases transforms marine piracy analysis by automating time-consuming tasks, enhancing speed, intelligence, and efficiency. The Phidata framework offers a powerful platform for processing unstructured data, enabling data scientists in maritime security to streamline workflows and extract actionable insights. By leveraging advanced AI agents, piracy data analysis becomes more efficient, supporting quicker, more accurate decision-making.

Looking ahead, multi-agent AI systems will evolve to play a more proactive role in maritime security. These systems will move beyond simple analysis, enabling real-time threat identification, automated alerts, and even autonomous decision-making. As AI agents advance, they will enhance predictive capabilities, adapt to emerging piracy trends, and reduce human intervention, ultimately creating a more responsive and secure maritime environment. This integration marks just the beginning of a new era in autonomous piracy analysis and prevention.

REFERENCES

- [1] A Phidata Framework Documentation. (2025). Retrieved from <https://docs.phidata.com/agents/storage>
- [2] Analyzing the Piracy from structured and unstructured data using SAS Text Miner <https://support.sas.com/resources/papers/proceedings15/3472-2015.pdf>
- [3] International Maritime Organization (IMO). (2024). Pirate Attacks Monthly Report. Retrieved from <https://wwwcdn.imo.org/localresources/en/OurWork/Security/Documents/Piracy%20monthly%20report%20November%202024.pdf>
- [4] Zhi Jing, Yongye Su, Yikun Han (2025). When Large Language Models Meet Vector Databases: A Survey. arXiv. Retrieved from <https://arxiv.org/abs/2402.01763>
- [5] Aniruddha Salve, Saba Attar, Mahesh Deshmukh, Sayali Shivpuje, Arnab Mitra Utsab (2025). A Collaborative Multi-Agent Approach to Retrieval-Augmented Generation Across Diverse Data. arXiv. Retrieved from <https://arxiv.org/abs/2412.05838>
- [6] Antony Seabra, Claudio Cavalcante, Joao Nepomuceno, Lucas Lago, Nicolaas Ruberg, Sergio Lifschitz. (2025). Dynamic Multi-Agent Orchestration and Retrieval for Multi-Source Question-Answer Systems using Large Language Models. arXiv. Retrieved from <https://arxiv.org/abs/2412.17964>
- [7] James Thorne, Majid Yazdani, Marzieh Saeidi, Fabrizio Silvestri, Sebastian Riedel, Alon Halevy(2025). Neural Databases. arXiv. Retrieved from <https://arxiv.org/abs/2010.06973>
- [8] Wikipedia contributors. (2025). LangChain. Wikipedia. Retrieved from <https://en.wikipedia.org/wiki/LangChain>
- [9] Phi Data framework <https://docs.phidata.com/introduction>