AI-Driven Evidence Synthesis: Leveraging RAG and multi-agentic approach to conduct disease landscape assessments

MSR16

Kasper M. Johannesen[1] Barinder Singh[2], Rajdeep Kaur[3],, Vicky Huang [4], Sven L. Klijn[4], Christian Boehler[5], Lisa Vaz[4], Matthew Walters[6], Nicola Waddell[2] Shubhram Pandey[3]

[1]Bristol Myers Squibb, Stockholm, Sweden, [2]Pharmacoevidence, London, UK, [3]Pharmacoevidence, Mohali, India,, [4]Bristol Myers Squibb, Princeton, NJ, USA, [5]Bristol Myers Squibb, Vienna, Austria, [6]Bristol Myers Squibb, Uxbridge, UK

Introduction

- Landscape assessments play a vital role in informing the early asset, business development and long-term strategic decision-making process, to optimize commercial success
- These assessments require integration of multiple data sources on disease background, treatment landscape and market access overview
- Manual landscape assessment can be time-consuming and error-prone due to changes in market dynamics and competitive settings
- Generative AI can automate operations like market landscape analysis and report generation using advanced techniques, resulting in simpler, faster, and more accurate workflows

Objectives

- The objective of this study was to develop an AI powered platform for conducting landscape assessments, including treatment guidelines, regulatory information and HTA assessments across different disease areas

Methods

- The platform integrated the dynamic Retrieval-Augmented Generation (RAG), multi agentic approach to analyze and summarize disease background, product information and access considerations across multiple countries
- It was built using Python microservices and a data processing pipeline. This interface used generative AI to automate and streamline the entire process of developing landscape assessment

Integration of data connectors

- The RAG based data pipeline enabled users to upload materials on disease background, current treatments and HTA relevant information for various disease areas
- The interface allowed the user to upload different file formats (PDF, Word, PowerPoint Deck and text files)
- Uploaded files are stored in an s3 bucket to extract the text, tabular and images

Integration of RAG pipeline

- The data processing pipeline process the data using two engines, a standardization engine which format the raw input into structured and consistent format and RAG engine which retrieve the relevant information and generate the content.
- The standardization engine converts the data into markdown format using Optical Character Recognition (OCR) as shown in Figure 1 (Phase-1)
- It extracts the structured elements such text in multiple columns, tables, images and graphs from the uploaded data and processed separately to ensure the high accuracy and maintain original clarity and structure
- RAG engine divided the markdown content into smaller chunks while keeping the context preserved and making it easier to retrieve
- The embedding were created from these markdown chunks and stored in the vector database for efficient semantic search
- It retrieve the most relevant chunks in response to the user query

Development of agents, tools, and prompt template

- LangChain-based framework used throughout the pipeline to provide evidence retrieval. It used customized prompt templates and agents with tools and memory to create relevant content from integrated data as shown in Figure 2
- Separate agents were designed to generate the content of different sections (disease overview, treatment landscape, HTA assessment overview etc.) of the landscape assessment
- This interface retrieves the data from the RAG as per the context provided by the user

Integration of LLM's and multiple agents in AI studio

The boto3_bedrock client was employed in the configuration to integrate the different LLM's:

- Amazon Titan : It was integrated with this interface which is responsible for embedding creation, it involves transforming data into high-dimensional vector representations.
- Model generated embeddings were stored in Postgres a database, which enabled efficient search and classification
- Claude Sonnet 3.5 v1 : This model was used in text analysis and generation. It played a key role in summarization, content creation, and natural language understanding. Streaming was set to 'True', to improve responsiveness for long outputs Claude Sonnet 3.5 v2 : This enhanced version of the Claude Sonnet model was used to perform image analysis and to generate the image description.
- The AI Studio (set of AI agents) as shown in Figure 3 was built on top of the RAG pipeline, enabling it to fetch relevant information from stored documents
- Different AI agents were designed to generate the content of different domains of landscape assessment (disease background, treatment landscape, guideline overview and HTA overview)

Figure 1. Standardization and RAG processing pipeline

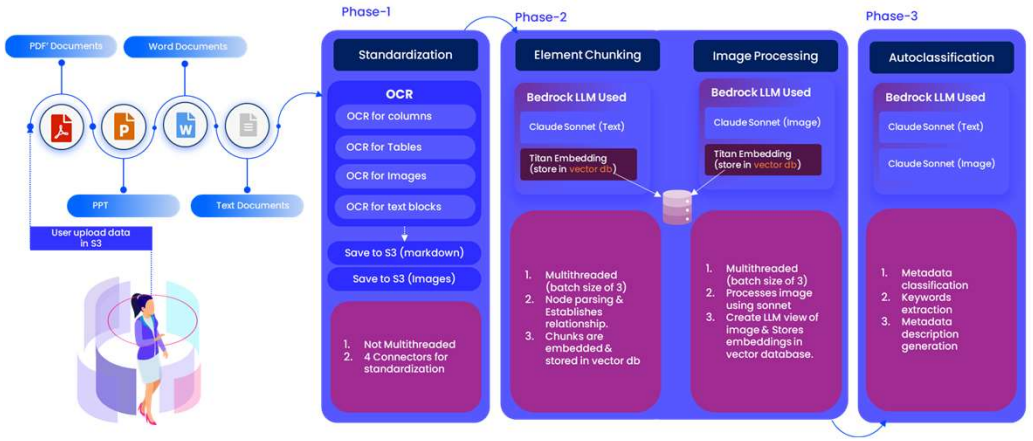

Figure 2. Landscape Assessment Framework: Multiple agents, memory, tools

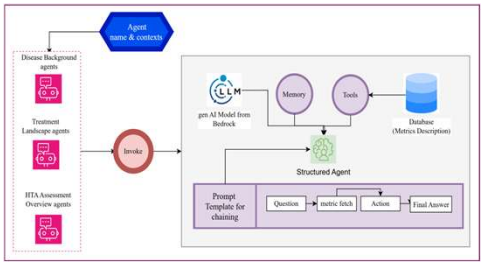

Figure 3. Landscape assessment agents framework

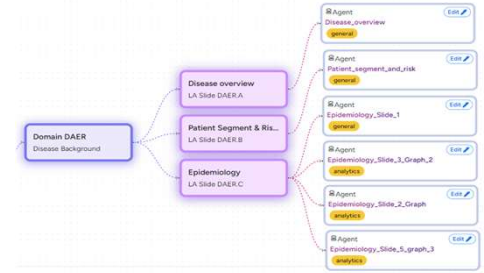

Analysis suite (multi-agentic response)

- Analysis suite was designed to present and map the results of multiple agents of the landscape assessment
- Every agent can extract the data from the vector database as per the context given by user and agent will generate the desired response
- Domain experts designed the context of different agents and selected the output type (General, Structured and Analytics)
  - General (paragraphs): This agent generates well structured markdown response in the form of text data
  - Structured (tables): This agent provides responses into clean HTML tables which further mapped into editable tables in Power Point template
  - Analytical (graphs/plots): This agent generates the python code to plot the graphs/plots using Plotly (Python library)

Template mapping (Power Point slide deck) and referencing

- Interface has functionality for the user to upload blank Power Point Template
- Different parser functions were built in Python to map the outputs in different formats (text, tables and graphs/plots)
- Each agent response is fetched from the S3 data buckets and then mapped into the uploaded template.
- Various Python libraries were used to format, structure and add subtitles for the content in the template
- After successful template mapping user can download the Power Point slide deck
- The platform included a referencing algorithm which refers the data and generate bibliographies accurately
- To validate the data extraction from the RAG pipeline, domain experts created 30 complex prompts on disease background, treatment landscape, guideline overview, HTA overview. product labels.
- The result of these prompts were evaluated manually
- This validation process was conducted by domain experts leveraging a 5-point Likert scale (Strongly Agree , Agree, Undecided, Disagree, Strongly Disagree) to evaluate the AI-generated output sets

Results

- The AI automated process is able to effectively synthesis data uploaded to the RAG engine, significantly reducing the time required to extract relevant data

Table 1. Overall distribution of the AI generated responses

Rating	Number of Responses	Percentage
Strongly Agree (5)	27	90.0%
Agree (4)	2	6.7%
Undecided (3)	1	3.3%
Disagree (2)	0	0.0%
Strongly Disagree (1)	0	0.0%
Total	30	

- Domain experts found that 27/30 of AI-generated output sets (90.0%) demonstrated strong alignment with human knowledge
- Two AI-generated outputs (2/30; 6.7%) showing relevant content but requirement refinement for improved clarity
- One AI generated output (1/30; 3.3%) was categorized as undecided, reflecting ambiguity where contextual information can be improved
- A sample prompt output is shown in Figure 4

Prompt

- Create a table listing different approved treatments for Myelofibrosis (MF). The table should include the following columns: Treatment Name; FDA/EMA Approved (Yes /No) along with date and month of approval; Indication of Use (provide separate indication for FDA and EMA); Key Differences in FDA and EMA Approval. Highlight any differences in indication, patient population, dosage, or restrictions between the FDA and EMA approvals. Ensure that all data is current and accurate.

Figure 4. Sample prompt result to validate the RAG pipeline

Amica.

Here's a table listing different approved treatments for Myelofibrosis (MF) with the requested information:

Treatment Name	FDA Approved	EMA Approved	FDA Indication	EMA Indication	Key Differences in FDA and EMA Approval
Ruxolitinib (Jakafi/Jakavi)	Yes, November 2011 [1]	Yes, August 2012 [2]	Treatment of intermediate or high-risk myelofibrosis, including primary MF, post-polycythemia vera MF and post-essential thrombocythemia MF [1]	Treatment of disease-related splenomegaly or symptoms in adult patients with primary MF, post-polycythemia vera MF or post-essential thrombocythemia MF [2]	EMA approval specifies "adult patients" while FDA does not have this age restriction explicitly stated
Fedratinib (Inrebic)	Yes, August 2019 [3]	Yes, February 2021 [4]	Treatment of adult patients with intermediate-2 or high-risk primary or secondary (post-polycythemia vera or post-essential thrombocythemia) MF [3]	Treatment of disease-related splenomegaly or symptoms in adult patients with primary or secondary MF who are JAK inhibitor naive or have been treated with ruxolitinib [4]	EMA approval includes patients who are JAK inhibitor naive or previously treated with ruxolitinib, while FDA approval does not specify prior treatment status
Pacritinib (Vonjo)	Yes, February 2022 [5]	No	Treatment of adults with intermediate or high-risk primary or secondary (post-polycythemia vera or post-essential thrombocythemia) MF with a platelet count below 50 × 10^9/L [5]	N/A	Not approved by EMA
Momelotinib (Ojjaara)	Yes, August 2023 [6]	Yes, January 2024 [7]	Treatment of intermediate or high-risk MF, including primary MF and secondary MF, in adults with anemia [6]	Treatment of disease-related splenomegaly or symptoms in adult patients with primary or secondary MF who are JAK inhibitor naive or have been treated with ruxolitinib [7]	FDA approval specifically mentions patients with anemia, while EMA approval includes JAK inhibitor naive patients or those previously treated with ruxolitinib

- In a test case on Myelofibrosis (MF), the AI studio generated a comprehensive and well structured PPR presentation detailing: Disease background; Treatment landscape; Guideline overview; and HTA overview.
- The AI studio enabled to user to adjust AI Agents' context and specify changes needed for the final PPT presentation generation
- Example of AI context and a slides from the generated presentation is shows in Figure 5 and Figure 6

Figure 5. AI Agent Context editor

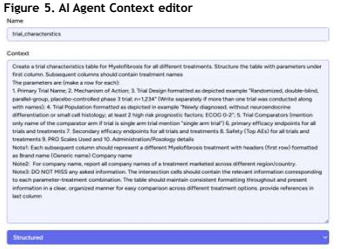

Figure 6. AI Studio with AI Agents supporting context modification functionality

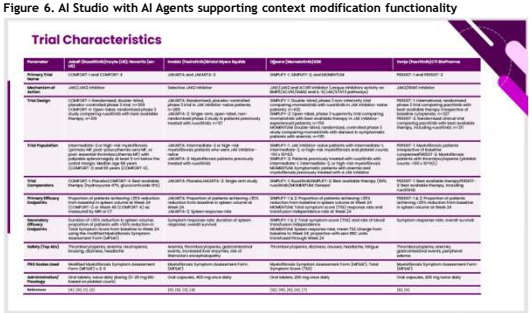

Conclusion

- The AI-powered platform demonstrated the capability to analyze and process large volumes of complex data from multiple sources like treatment guidelines, HTAs and regulatory documents, with a high degree of accuracy.
- The automation of landscape assessments, using data pre-processing engines (standardization engine and RAG engine) and other generative AI approaches, presents an approach which offers the potential for scalability. Additionally, generalization to other data synthesis and alternative use cases is possible and will require validation in future research.

References

[1] U.S. Food and Drug Administration. FDA approves first drug to treat a rare bone marrow disease. 2011.
[2] European Medicines Agency. Jakavi: EPAR - Product Information. 2012.
[3] U.S. Food and Drug Administration. FDA approves fedratinib for myelofibrosis. 2019.
[4] European Medicines Agency. Inrebic: EPAR - Product Information. 2021.
[5] U.S. Food and Drug Administration. FDA Approves Drug for Adults with Rare Form of Bone Marrow Disorder. 2022.
[6] U.S. Food and Drug Administration. FDA Approves Momelotinib for Myelofibrosis with Anemia. 2023.
[7] European Medicines Agency. Ojjaara: EPAR - Product Information. 2024.
[8] Gerds A, et al. NCCN clinical practice guidelines in oncology: myeloproliferative neoplasms. National Comprehensive Cancer Network (NCCN) 2024
[9] Committee N, et al. NICE guidance on fedratinib for myelofibrosis-related splenomegaly and symptoms. National Institute for Health and Care Excellence (NICE) 2021
[10] Committee N, et al. NICE guidance: momelotinib for treating myelofibrosis-related splenomegaly or symptoms. National Institute for Health and Care Excellence (NICE). 2024
[11] Jakafi® (ruxolitinib). Wilmington, DE: Incyte Corporation; 2023
[12] Leonard B, et al. Patient and clinician input on momelotinib for myelofibrosis treatment. CADTH. 2023