

# Comprehensive Insights into Global Mineral Commodities: Analysis, Visualization and Intelligent Assistance

Trupti Mohanty<sup>1†</sup>, Hasan M. Sayeed<sup>1†</sup>, Chitrasen Mohanty<sup>2</sup>, Taylor D. Sparks<sup>1\*</sup>

<sup>1</sup>*Department of Materials Science & Engineering, University of Utah, Salt Lake City, UT-84112, USA.*

<sup>2</sup>*Department of Biostatistics & Medical Informatics, University of Wisconsin Madison, Madison, WI-53726, USA.*

† Authors contributed equally.

\***Correspondence to:** Prof. Taylor D. Sparks, Department of Materials Science & Engineering, University of Utah, Salt Lake City, UT-84112, USA, E-mail: [sparks@eng.utah.edu](mailto:sparks@eng.utah.edu)

## Abstract

With the growing emphasis on sustainability, criticality, and availability in materials research, our study introduces a comprehensive data analytics platform to provide country-specific insights into global elemental production and reserves. Utilizing data from the United States Geological Survey (USGS), our web application incorporates the Herfindahl-Hirschman Index (HHI) to assess market concentration, identifying potential risks and opportunities related to resource availability. The platform features an AI assistant powered by a Retrieval-Augmented Generation (RAG) system, leveraging the past ten years of USGS mineral commodities summaries. This system employs an open-source large language model (LLM) to enable users to query various aspects of raw materials, including reserves, production, market share, usage, price, substitutes, recycling, and more. By retrieving relevant documents and generating accurate, comprehensive responses, our tool addresses a crucial gap in publicly available resources, offering a unified application for detailed material analysis. This platform provides valuable support for material scientists in assessing sustainability, criticality, and market risks, thereby aiding in the development of new materials. Website: <https://mineral-ai.net>

**Keywords:** United States Geological Survey (USGS), Herfindahl-Hirschman Index (HHI), Retrieval-Augmented Generation (RAG), Large Language Model (LLM)

## 1. Introduction

Understanding the availability and risks associated with raw materials is crucial for numerous fields, particularly those focused on materials science, industrial ecology, and sustainability. As global demand for various elements increases, monitoring and analyzing their production and reserves becomes increasingly important. This knowledge helps researchers, policymakers, and industry stakeholders make informed decisions, ensuring the sustainability and security of supply chains.

Raw materials are fundamental to technological advancement and economic growth. Elements such as lithium, cobalt, and rare earths are critical for developing low-carbon technologies and renewable energy systems. These materials are essential for manufacturing batteries, wind turbines, and electric vehicles, which play a pivotal role in reducing greenhouse gas emissions and combating climate change [1]. Any disruption in the supply of these materials can have significant ramifications, affecting everything from manufacturing processes to geopolitical stability. Consequently, understanding the global distribution and market concentration of these elements is essential for ensuring a stable supply chain.

To evaluate market concentration, one of the primary metrics used is the Herfindahl-Hirschman Index (HHI). The HHI quantifies the market share distribution among countries, providing insights into the level of competition and market dominance in the global landscape [2]. A high HHI value suggests a market dominated by a few producers, which increases the risk of supply disruption and price volatility. Conversely, a low HHI indicates a more competitive market with multiple producers, thereby reducing these risks. This metric has been applied to study the market concentration of the materials associated with thermoelectric and battery technologies [3, 4], revealing the potential risks associated with high HHI values. These studies underscore the vulnerabilities, such as supply risks and price volatility, that arise when a market is dominated by a few producers, highlighting the crucial need for selecting materials with lower HHI values to ensure a stable and reliable supply chain.

Further emphasizing the significance of HHI, several studies have highlighted the importance of understanding raw material availability and its associated risks. For instance, recent studies on lithium supply highlighted the significant challenges in meeting future demand due to geopolitical and environmental constraints [5-7]. Another study analyzed the supply chain of cobalt, emphasizing the risks posed by its concentrated production in politically unstable regions [8]. Research on rare earth elements has underscored the criticality of these materials for green technologies and the potential supply risks due to limited global production [9]. Additionally, Tehrani et al. (2017) explored the balance between mechanical properties and sustainability in the search for superhard materials, highlighting the necessity for a stable supply of raw materials to support innovative manufacturing processes [10].

Despite the extensive research and insights provided by these studies, there remains a significant gap in the availability of a comprehensive, publicly accessible resource for visualizing and exploring data on the reserves and production of various elements across countries and over time. The United States Geological Survey (USGS) Minerals Yearbook and Mineral Commodity Summaries provide extensive data [11], but these resources are often difficult to analyze in detail and lack advanced querying features for extracting specific information and performing complex queries.

To bridge this gap, our study introduces a comprehensive data analytics platform that provides country-specific insights into global elemental production and reserves, along with AI-assisted querying for raw materials. At the core of this platform are Large Language Models (LLMs), which have revolutionized natural language processing by showcasing impressive capabilities in understanding and generating human-like text. These models, trained on vast amounts of data, can perform tasks such as translation, summarization, and question-answering. However, LLMs also have limitations, including generating plausible-sounding but incorrect answers (hallucinations), struggling with outdated knowledge, and having opaque and untraceable reasoning processes [12]. To address these issues, we employ a Retrieval-Augmented Generation (RAG) system. RAG enhances LLMs by integrating external knowledge sources into the generation process. It retrieves relevant documents based on the user's query and combines this information with the LLM's intrinsic knowledge to produce more accurate and contextually relevant responses. The typical RAG workflow involves indexing documents, retrieving the top relevant chunks using semantic similarity, and generating a final answer using both the query and the retrieved information. This

approach significantly reduces hallucinations and ensures that the generated content is up-to-date and traceable [13, 14].

To meet the needs of researchers, policymakers, and industry stakeholders, we have developed a web application leveraging USGS data to offer detailed insights into global mineral commodities. This platform incorporates the HHI to evaluate market concentration and potential risks. An AI assistant, powered by a RAG system using an open-source LLM, facilitates queries on various aspects of raw materials. By retrieving relevant documents and generating comprehensive responses, this tool provides a unified platform for assessing sustainability, criticality, and market risks in material science.

## 2. Methods

### 2.1 Data Collection and Pre-processing

The data for global elemental production and reserves from 2016 to 2023 was sourced from the USGS Minerals Yearbook and Mineral Commodity Summaries. These comprehensive resources provide detailed, country-wise data on production and reserves. The pre-processing phase involved cleaning and organizing the data to ensure consistency and accuracy by handling missing values, standardizing units of measurement and saving in CSV files.

We focused on various elements and calculated their respective market shares by country for each year, followed by the calculation of the HHI. The HHI is a widely recognized measure of market concentration, defined as:

$$\text{HHI} = \sum_i^N S_i^2$$

Where,  $N$  is the total number of countries involved in the world production or reserve for a given metal/mineral and  $S_i$  is the percent market share of the country  $i$ . The HHI ranges from 0 to 10,000, where  $\text{HHI} = 0$  indicates a market controlled by a large number of countries with nearly equal shares, and  $\text{HHI} = 10,000$  indicates a market controlled by a single country. According to the U.S. Department of Justice and Federal Trade Commission, the concentration of the market can be categorized based on HHI values as follows: an HHI value of less than 1500 indicates an unconcentrated market, an HHI value between 1500 and 2500 indicates a moderately concentrated market, and an HHI value greater than 2500 indicates a highly concentrated market. These standard metrics were utilized in our study to visualize the market risk for different elements [2].

### 2.2 Visualization

We developed a web-based application to facilitate the interactive exploration of material availability and market risks across various countries globally. The technical architecture of the application was built using HTML (HyperText Markup Language), CSS (Cascading Style Sheets), and JavaScript, and it is hosted on a Python based web server [15]. The front-end interface was designed as a single-page layout using the Bootstrap [16] and Plotly frameworks [17]. The periodic table allows users to select multiple elements for exploration. Selected elements are highlighted,

providing a visual cue for user selections. Users can customize their queries by specifying desired years and choosing between reserves or production data.

The system dynamically retrieves data from the CSV files and presents it in various graphical formats. Color-coded world maps illustrate the geographical distribution of market shares for reserve and production. Users can hover over countries to view specific market shares in percentages and absolute values in tons. Line charts track HHI values over time, categorized by market concentration levels (high, moderate, or unconcentrated), providing a temporal perspective on market dynamics. Bar charts compare reserve and production market shares of various elements across different countries and years. Additionally, users can download data in CSV format for independent analysis, further research, or archival purposes. This feature enhances the application's utility for academic research and industry stakeholders. Figure 1. provides a comprehensive visualization of the selected elements and the chosen year. The world map illustrates the production and reserve market share for different countries for the highlighted element lithium in 2023. Additionally, it displays the HHI of all selected elements and a market share bar graph for lithium.

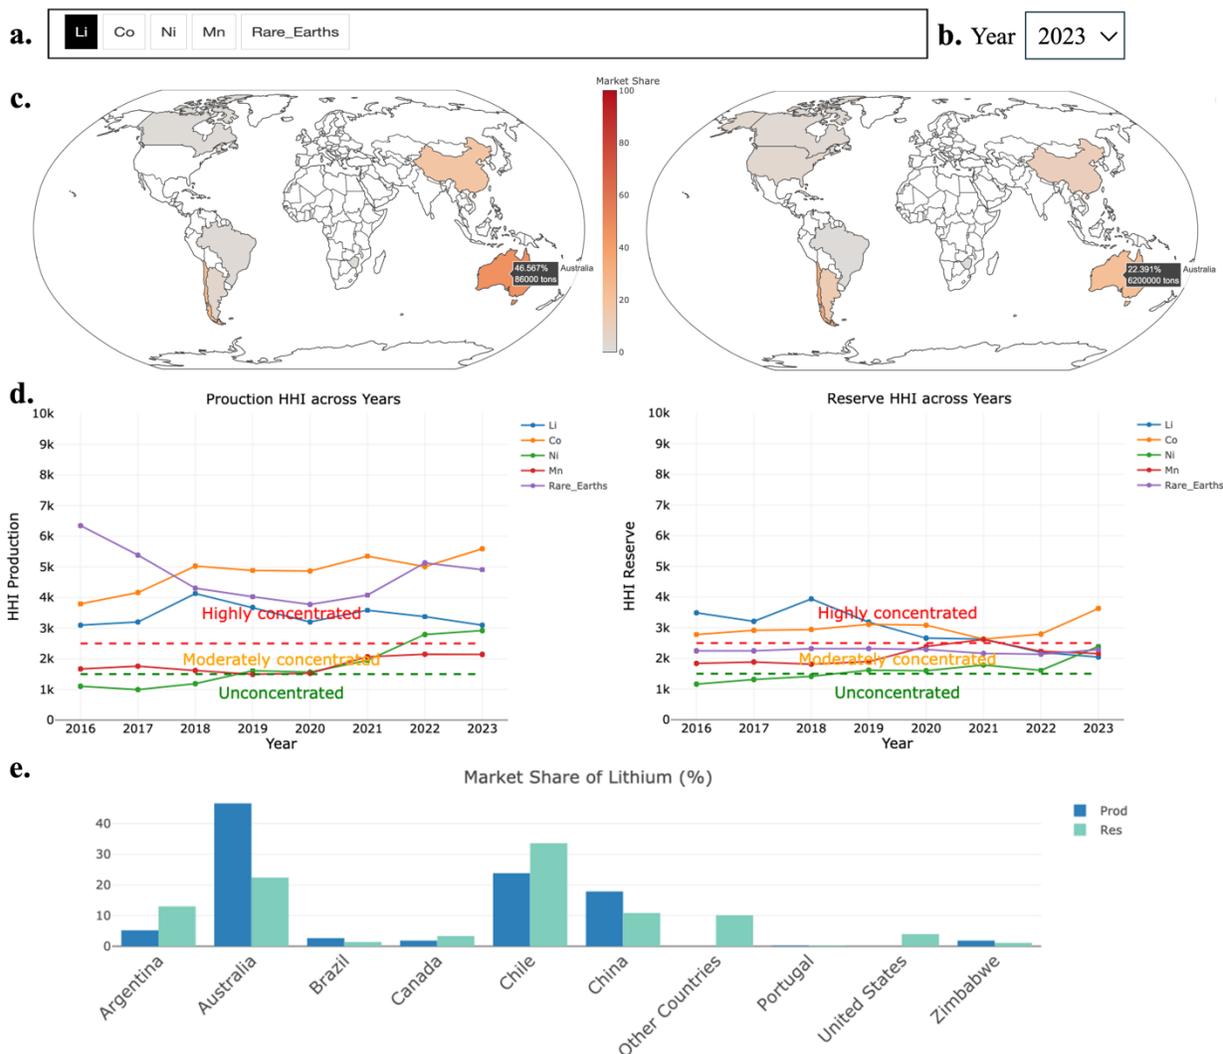


Figure 1. illustrates a comprehensive visualization of the selected elements and the chosen year. a. The element

selection panel shows the selected elements. **b.** The year selection dropdown enables users to specify the desired year for querying data, allowing for temporal analysis of production and reserve. **c.** Color-coded world maps illustrating the geographical distribution of market shares for production (left) and reserve (right) of the highlighted element. Users can hover over countries to view specific market shares in percentage and tons as shown for Australia. **d.** Line charts displaying the HHI values over time for production (left) and reserves (right) of the selected elements categorizing market concentration levels into high, moderate, and unconcentrated. **e.** Bar charts comparing the market shares of reserves and production for the selected element across different countries for the chosen year.

### 2.3 AI Assistant (MatAssist) for Querying Mineral Commodities

Central to our platform is an AI assistant, MatAssist, powered by a Retrieval-Augmented Generation (RAG) system, specifically designed to process element-wise PDFs of mineral commodity summaries from the USGS for the past decade. The overall framework of our platform was developed using LangChain, an open-source software library designed for the rapid development of LLM applications [18, 19]. Figure 2. shows the overall workflow of the AI assistant.

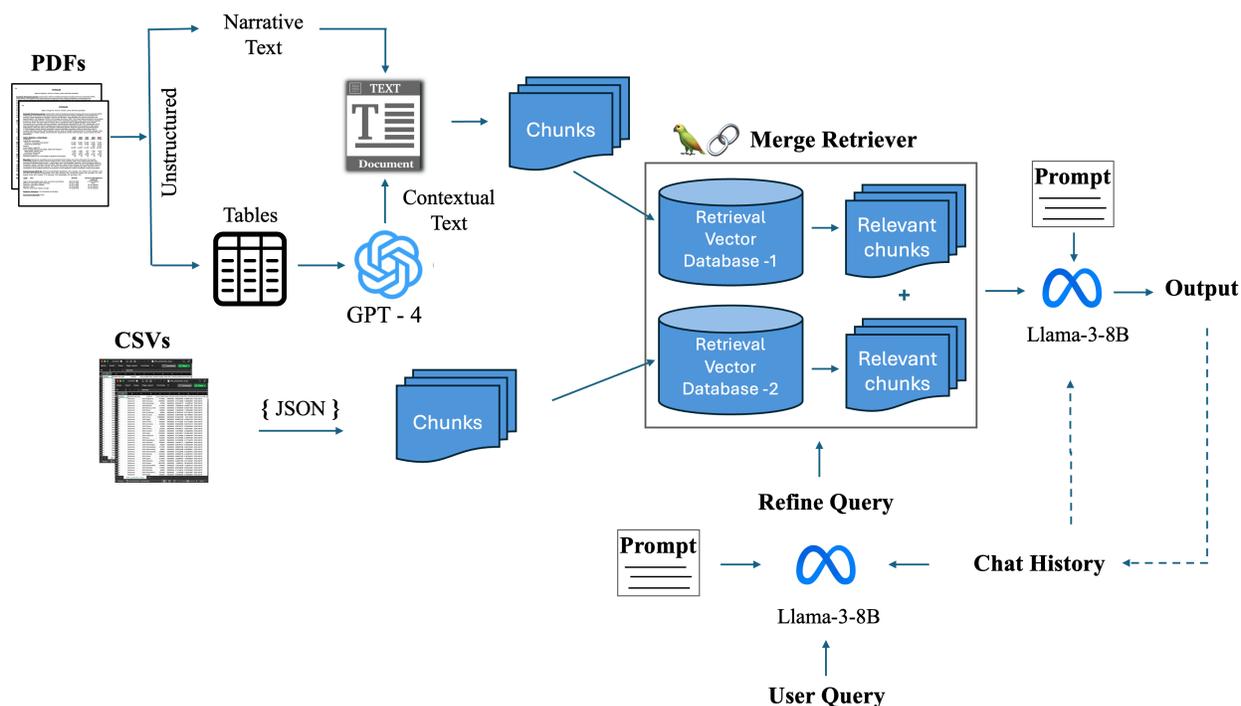


Figure 2. shows the workflow of the AI assistant. Text is extracted from PDFs and CSVs, with tables converted to contextual text using GPT-4. The text is split into chunks and encoded into vector databases. Relevant chunks specific to user queries are retrieved, and responses are generated by the LLM using these chunks and chat history.

The PDF documents contain critical information on materials' domestic production, usage, imports, exports, prices, recycling, substitutes, events, trends, issues, world mine production, and reserves. The unstructured nature of PDFs poses significant challenges for extracting accurate data, particularly from tables. LLMs often struggle to accurately process and answer queries involving tabular data within these PDFs. Additionally, inconsistent formatting across PDFs and the inability

to visually identify and classify document elements can result in incomplete or incorrect information retrieval.

To address these challenges, we utilized the Unstructured API [20], which employs document layout detection through the YOLOX model [21]. This model identifies and draws bounding boxes around various document elements, including titles, narrative texts, and tables. The text within the table elements is processed using the GPT-4 [22] large language model to convert them into plain contextual text based on user prompt. The user prompt, example table, and corresponding output are detailed in Table TS1 of the supplementary material. The narrative text, along with the contextual text from GPT-4, forms a comprehensive text document for each PDF. This document is then split into small chunks and encoded into vector representations using the Hugging Face UAE-Large-V1 embedding model [23], chosen for its balance of memory efficiency and embedding accuracy [24]. The encoded vectors are stored in a vector database using ChromaDB, an open-source vector store designed to manage high-dimensional vector embeddings, allowing efficient storage, indexing, and retrieval functionality [25].

CSV documents containing year-wise market share and HHI data for reserves and production are converted into JSON format and are processed similarly, split, and stored in a separate vector database. Upon receiving a user query, the system reformats the query to ensure contextual relevance using the question formation prompt and the chat history detailed in TS2 of the supplementary material. This process is facilitated by the LLMChain module from LangChain, utilizing the open-source LLM “Meta-Llama-3-8B-Instruct” [26] available in the HuggingFaceHub library. The reformatted query is then processed by both vector retrievers retrieving the top K chunks based on the similarity scores between the query vector and the vectors of chunks. The LangChain MergeRetriever combines results from these retrievers into a single list, enhancing the accuracy of document retrieval by increasing variability and reducing the risk of bias.

The final response is generated by the LLM “Meta-Llama-3-8B-Instruct,” considering the merged context and the conversation history. This process ensures that the response is informative and interactive, effectively handling formatting and source document inclusion. The combination of LangChain for workflow orchestration, ChromaDB for efficient vector management, Unstructured for handling tables, GPT-4 for contextual information extraction from tables, and an open-source LLM for coherent response generation creates a robust and reliable system to address the complexities of extracting and interpreting data from unstructured PDF documents.

### **3. Results and Discussion**

#### **3.1 Geographical Distribution and HHI Trends Over Time**

Our platform provides visualizations of the geographical distribution of mineral production and reserves, market concentration, and HHI trends over time. Figure 1 shows data for key elements such as lithium (Li), cobalt (Co), nickel (Ni), manganese (Mn), and rare earths. For example, the world map in Figure 1. c highlights the primary producers and reserves of the highlighted element lithium in 2023, with Australia, Chile, and China being the leading producers. Australia and Chile also hold significant reserves. The HHI values in Figure 1. d reveal market concentration for these

elements, with cobalt and rare earths showing high production HHI values, indicating significant market control by a few countries. Analyzing HHI trends over time helps us understand the temporal dynamics of mineral markets, capturing the impacts of geopolitical events and data availability for certain countries. These insights are essential for assessing potential supply risks and market dynamics. For a detailed discussion and specific examples related to energy materials, we have provided an in-depth analysis in Section 4.

### 3.2 RAG Evaluation

To evaluate the quality of responses generated by MatAssist, metrics such as ROUGE and BERTScore are employed, focusing on aspects like factual correctness, readability, and user satisfaction [27]. ROUGE (Recall-Oriented Understudy for Gisting Evaluation) measures the overlap of n-grams between generated and reference texts [28], while BERTScore uses contextual embeddings from pre-trained transformers to evaluate semantic similarity [29].

The evaluation involved 30 user queries covering various topics such as market share, HHI, imports, exports, prices, events, trends, domestic production, usage, recycling, and substitutes. For each query, MatAssist's responses were compared with human-generated ground truths based on the original documents to assess accuracy and relevance. The details of these queries, MatAssist's responses, and the corresponding ground truths are provided in Table TS3 of the supplementary material. These ground truths were used as benchmarks to evaluate the performance of the RAG system. The scores obtained are shown in Figure 3. These results indicate that the system performs well in maintaining the contextual integrity and relevance of the responses. The high ROUGE recall values demonstrate the system's capability to capture a broad range of relevant information, ensuring that critical content is included in the responses. The high BERTScore values highlight the system's strength in capturing the semantic meaning of the queries and providing responses that are not only factually correct but also contextually appropriate [27, 30]. This comprehensive performance demonstrates the utility of the RAG system in generating accurate and informative outputs for complex queries related to mineral commodities, supporting informed decision-making for researchers, policymakers, and industry stakeholders.

While our platform provides comprehensive insights and an effective AI assistant for querying mineral commodity data, there are certain limitations to consider. One notable limitation is that the AI assistant may sometimes merge information from different years when the same data is available across multiple years, potentially leading to less precise responses. To mitigate these issues, it is recommended to specify the year when making queries to ensure more accurate and contextually relevant responses. Additionally, for queries that require extensive information, the response might be incorrect if the relevant data cannot be accommodated within the number of chunks processed by the AI. Our element-specific data represented in bar graphs complements this by providing an intuitive way to explore and verify the data.

Furthermore, not all elements in the periodic table have complete data available on the USGS site. For example, rare earth elements are grouped together which includes scandium (Sc) and yttrium (Y) combined with lanthanides. Similarly, platinum group metals, such as iridium (Ir), osmium (Os), palladium (Pd), platinum (Pt), rhodium (Rh), and ruthenium (Ru) are often reported together.

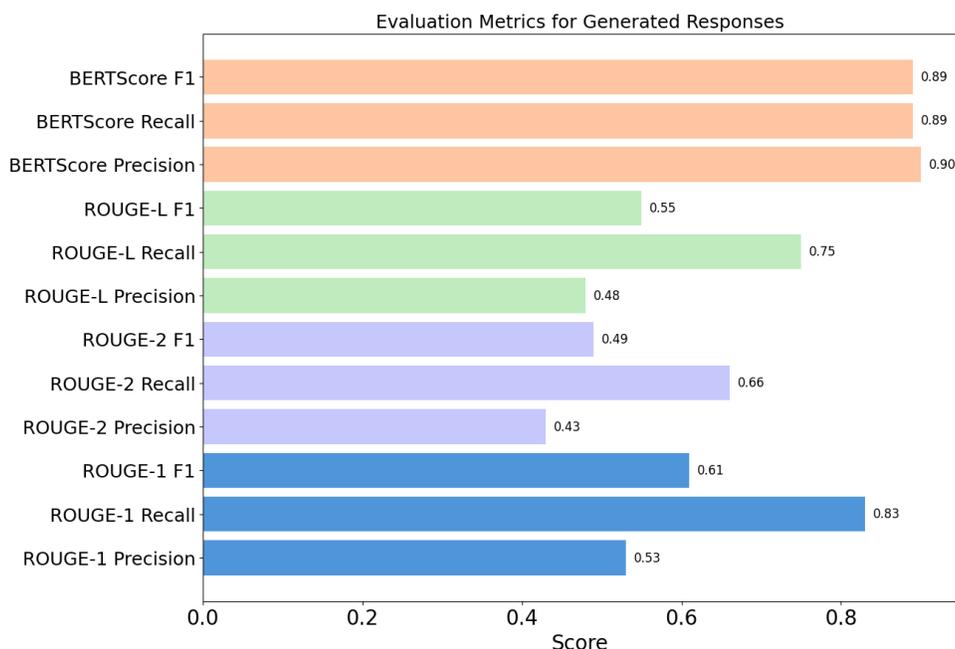


Figure 3. Evaluation metrics for the generated responses by MatAssist compared to ground truth data

#### 4. Use Case: Energy Materials

To illustrate the practical application of our platform, we demonstrate the analysis of critical elements essential for the development of batteries, wind turbines, and electric vehicles [1]. The primary elements analyzed include lithium, cobalt, nickel, manganese, and rare earth elements.

##### 4.1 Analysis of Market Share and HHI for Energy Materials

The provided figures (Figure 4) illustrate the market share percentages for the production and reserves of critical energy materials lithium, cobalt, nickel, manganese, and rare earth elements across various countries. Figure 1.d represents their HHI trends over time of these elements.

**Lithium:** The lithium market is predominantly controlled by Australia and Chile, with Australia having the largest share in production, while Chile holds significant reserves. This concentration underscores the strategic importance of these two countries in the global lithium supply chain. China's presence is also notable, indicating its growing influence in the lithium market. For lithium, the production HHI in 2023 was 3096, indicating a highly concentrated market, while the reserve HHI was 2048, suggesting moderate concentration.

**Cobalt:** Cobalt production is highly concentrated in the Congo (Kinshasa), which accounts for over 70% of global production. This reliance on Congo (Kinshasa) poses significant supply risks due to political and economic instability. Other notable producers include Indonesia and Russia, but their contributions are considerably smaller. Cobalt displayed higher concentration levels

compared to lithium, with a production HHI of 5591 and a reserve HHI of 3626, indicating significant market control by a few countries.

**Nickel:** Indonesia is the leading producer of nickel, with a substantial share of global production, followed by the Philippines. Indonesia also holds significant nickel reserves. The diversified production landscape for nickel indicates a relatively lower supply risk compared to cobalt. Figure 1.d (left) shows the production HHI trends over the years, highlighting that nickel's HHI has been increasing since 2016. Year-wise production analysis reveals that this increase is due to Indonesia's growing production market share.

**Manganese:** Australia and South Africa are key players in the manganese market, with significant shares in both production and reserves. The presence of multiple producing countries, including Brazil, China, Gabon, Ghana, and India, indicates a well-distributed supply chain, reducing the overall risk of supply disruptions. The production HHI for manganese is 2142, which is relatively lower compared to other critical elements.

**Rare Earth Elements:** China dominates the rare earth elements market, with the largest shares in both production and reserves. This concentration highlights the critical dependency on China for these essential materials used in various high-tech and renewable energy applications. Other notable reserves are found in Brazil and Vietnam. The production market for rare earth elements is highly concentrated with an HHI of 4909, and the reserve market is moderately concentrated with an HHI of 2288.

The analysis of these figures underscores the critical role of specific countries in the supply chains of essential energy materials. The concentration of production and reserves in a few key regions, particularly for cobalt and rare earth elements, emphasizes the need for strategic planning and international cooperation to ensure stable and sustainable supply chains.

## 4.2 AI Assistant Interaction for Energy Materials

MatAssist provides valuable insights and detailed responses to queries related to the production, price trends, substitutes, imports, exports, and recycling resources of critical energy materials. Here are few example interactions showcasing its capabilities presented in Table 1:

Table 1. Few examples of interactions with MatAssist for energy materials

User Query → What was the price trend for lithium from 2021 to 2023?
MatAssist → The historical price trend of lithium from 2021 to 2023 is as follows: <ul style="list-style-type: none"><li>• In 2021, the price of battery-grade lithium carbonate was \$12,600 per metric ton.</li><li>• In 2022, the price of battery-grade lithium carbonate significantly increased to \$68,100 per metric ton.</li><li>• In 2023, the price of battery-grade lithium carbonate decreased to \$46,000 per metric ton.</li></ul> This trend indicates a significant increase in lithium prices from 2021 to 2022, followed by a decrease in 2023.
User Query → What are the primary substitutes for Cobalt in lithium-ion battery manufacturing?

---

MatAssist → The primary substitutes for Cobalt in battery manufacturing, specifically in lithium-ion batteries, are:

Iron and Phosphorus: These elements can be used to create cobalt-free substitutes for lithium-ion batteries.

Note that the suitability of these substitutes may vary depending on the specific application and the desired performance characteristics.

---

User Query → What is the production trend of rare earths in China from 2020 to 2023?

---

MatAssist → The production trend of rare earths in China from 2020 to 2023 is as follows:

- 2020: 140,000 tons
  - 2021: 168,000 tons
  - 2022: 210,000 tons
  - 2023: 240,000 tons
- 

User Query → What are the recycling resources available for Manganese?

---

MatAssist → Recycling resources available for manganese in the United States are limited. Manganese was recycled incidentally as a constituent of ferrous and nonferrous scrap; however, scrap recovery specifically for manganese was negligible. Manganese is recovered along with iron from steel slag.

---

User Query → Briefly explain about World Resources of Ni in 2023.

---

MatAssist → Here's a brief overview of the world resources of Nickel (Ni) in 2023:

- Globally, nickel resources have been estimated to contain more than 350 million tons of nickel, with 54% in laterites and 35% in magmatic sulfide deposits. Hydrothermal systems, seafloor manganese crusts and nodules, and miscellaneous resources contain 10% and 1%, respectively.
  - Identified land-based resources averaging approximately 0.5% nickel or greater contain at least 300 million tons of nickel, with about 60% in laterites and 40% in sulfide deposits.
  - Extensive nickel resources are also found in manganese crusts and nodules on the ocean floor.
- 

MatAssist effectively provides detailed and insightful responses to queries related to energy materials. This comprehensive approach supports informed decision-making for researchers, policymakers, and industry stakeholders, facilitating the transition to renewable energy systems and technologies.

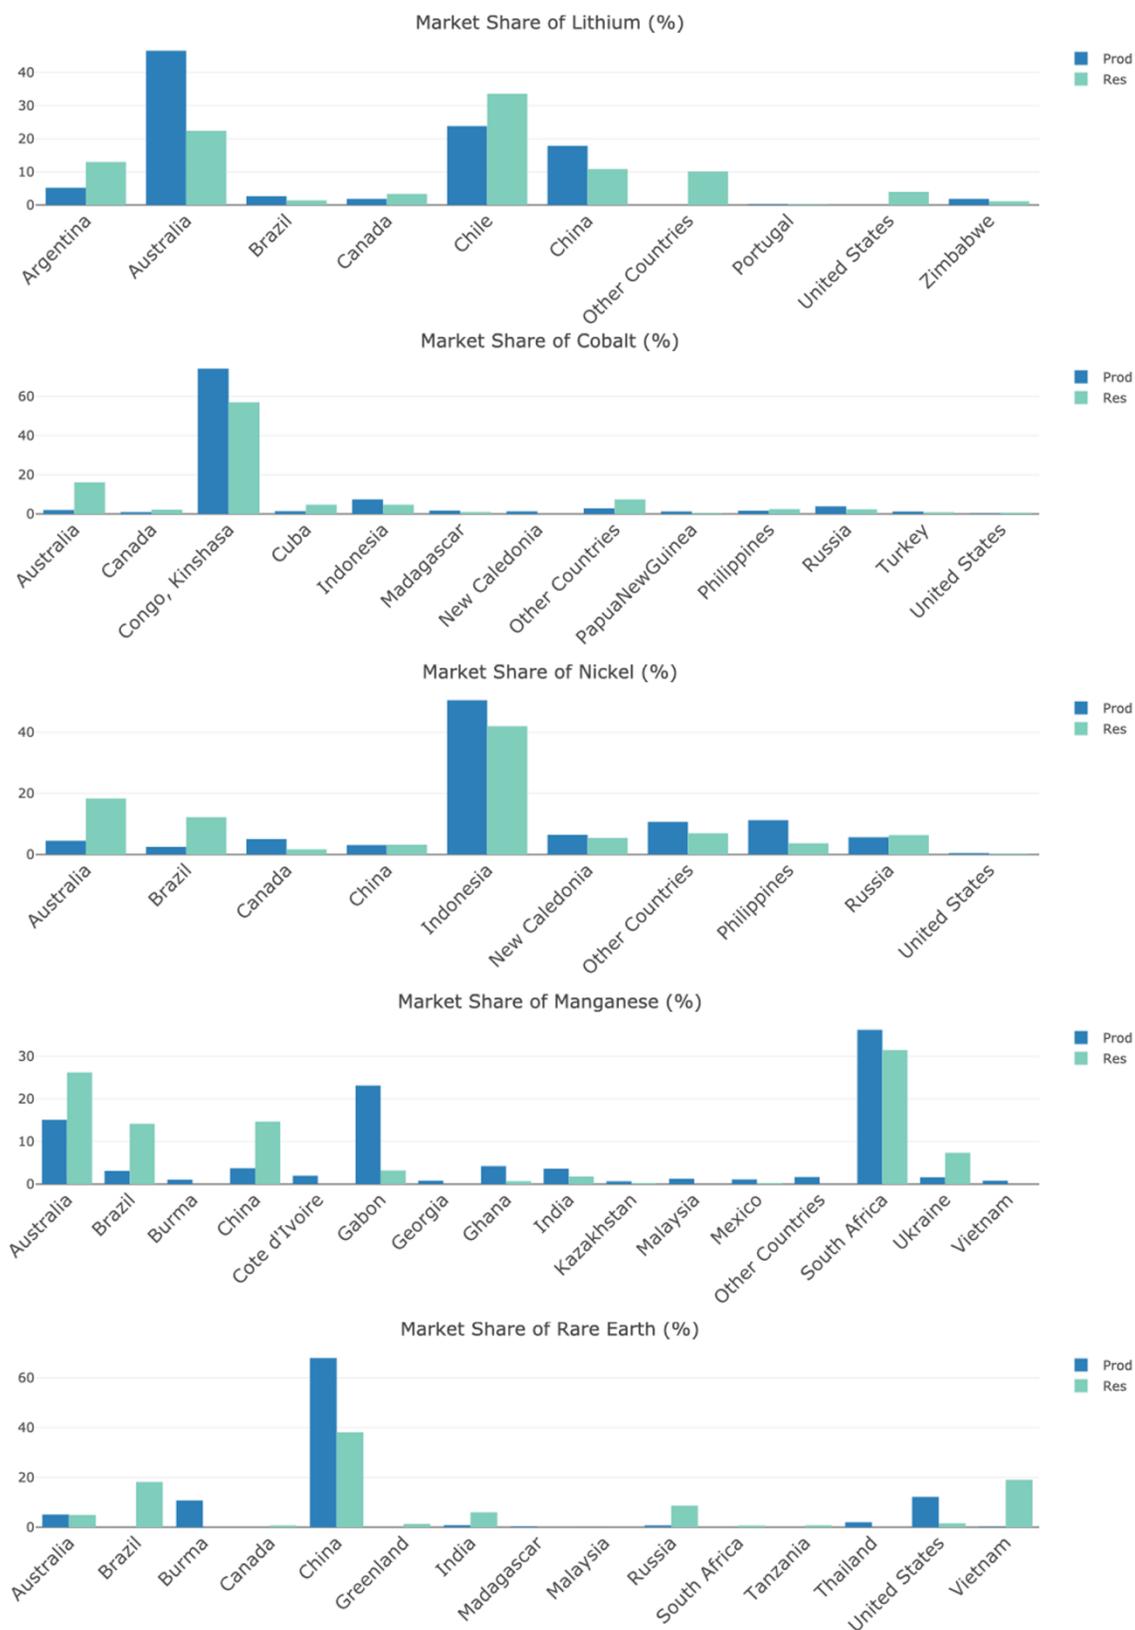


Figure 4. Market share production and reserve for the energy materials

## 5. Conclusions

The comprehensive data analytics platform developed in this study offers significant insights into global mineral commodities, with a focus on market concentration and supply risks. The integration of MatAssist, an AI assistant powered by a RAG system, enhances the platform's ability to provide detailed, accurate responses to complex queries. This supports informed decision-making for researchers, policymakers, and industry stakeholders. This tool is particularly valuable for assessing the sustainability and criticality of materials essential for developing new technologies and high-demand applications. By offering an integrated source of data and advanced analytics, we anticipate that this platform will significantly contribute to research efforts towards sustainability.

## Supplementary Material

The supplementary material includes prompts used for GPT-4 to derive contextual text from the tables, as well as prompts for MatAssist to reformat queries and generate responses. It contains the user queries, sample responses, and ground truths for all 30 queries used in the RAG evaluation metrics.

## Acknowledgments

The authors gratefully acknowledge funding from the Army Research Office Materials Design program under contract number #W911NF-23-1-0333.

## Author Declarations

Conflict of Interest: The authors have no conflicts to disclose.

## Data Availability

The source codes, datasets, and algorithms used in this study are available on GitHub at <https://github.com/truptimohanty/mineral-ai>

## References

1. Junne, Tobias, et al. "Critical materials in global low-carbon energy scenarios: The case for neodymium, dysprosium, lithium, and cobalt." *Energy* 211 (2020): 118532.
2. Theler, Brennan, Steven K. Kauwe, and Taylor D. Sparks. "Materials Abundance, Price, and Availability Data from the Years 1998 to 2015." *Integrating Materials and Manufacturing Innovation* 9 (2020): 144-150.
3. Gaultois, Michael W., et al. "Data-driven review of thermoelectric materials: performance and resource considerations." *Chemistry of Materials* 25.15 (2013): 2911-2920.

4. Ghadbeigi, Leila, et al. "Performance and resource considerations of Li-ion battery electrode materials." *Energy & Environmental Science* 8.6 (2015): 1640-1650.
5. Olivetti, Elsa A., et al. "Lithium-ion battery supply chain considerations: analysis of potential bottlenecks in critical metals." *Joule* 1.2 (2017): 229-243.
6. Altiparmak, Suleyman Orhun. "China and lithium geopolitics in a changing global market." *Chinese Political Science Review* 8.3 (2023): 487-506.
7. Alessia, Amato, et al. "Challenges for sustainable lithium supply: A critical review." *Journal of Cleaner Production* 300 (2021): 126954.
8. Liu, Wei, et al. "Resilience assessment of the cobalt supply chain in China under the impact of electric vehicles and geopolitical supply risks." *Resources Policy* 80 (2023): 103183.
9. Mancheri, Nabeel A., et al. "Effect of Chinese policies on rare earth supply chain resilience." *Resources, Conservation and Recycling* 142 (2019): 101-112.
10. Mansouri Tehrani, Aria, et al. "Balancing mechanical properties and sustainability in the search for superhard materials." *Integrating materials and manufacturing innovation* 6 (2017): 1-8.
11. United States Geological Survey (USGS) Minerals Yearbook and Mineral Commodity Summaries <https://www.usgs.gov/centers/national-minerals-information-center/minerals-yearbook-metals-and-minerals>
12. Tonmoy, S. M., et al. "A comprehensive survey of hallucination mitigation techniques in large language models." arXiv preprint arXiv:2401.01313 (2024).
13. Lewis, Patrick, et al. "Retrieval-augmented generation for knowledge-intensive nlp tasks." *Advances in Neural Information Processing Systems* 33 (2020): 9459-9474.
14. Gao, Yunfan, et al. "Retrieval-augmented generation for large language models: A survey." arXiv preprint arXiv:2312.10997 (2023).
15. Web technology for developers, <https://developer.mozilla.org/en-US/docs/Web>
16. Bootstrap: The world's most popular framework for building responsive, mobile-first sites. Retrieved from <https://getbootstrap.com/>
17. Plotly Technologies Inc. "Plotly: Collaborative Data Science." 2015, plotly.com.
18. Topsakal, Oguzhan, and Tahir Cetin Akinci. "Creating large language model applications utilizing langchain: A primer on developing llm apps fast." *International Conference on Applied Engineering and Natural Sciences*. Vol. 1. No. 1. 2023.
19. LangChain Documentation, <https://python.langchain.com/v0.2/docs/introduction/>
20. Unstructured Documentation, <https://docs.unstructured.io/welcome>
21. Ge, Zheng, et al. "Yolox: Exceeding yolo series in 2021." arXiv preprint arXiv:2107.08430 (2021).

22. OpenAI. "GPT-4." OpenAI, 2023, [www.openai.com/research/gpt-4](http://www.openai.com/research/gpt-4)
23. Universal AnGLE Embedding, <https://huggingface.co/WhereIsAI/UAE-Large-V1>
24. Massive Text Embedding Benchmark Leaderboard, <https://huggingface.co/spaces/mteb/leaderboard>
25. ChromDB Documentation, <https://docs.trychroma.com/getting-started>
26. "Meta-Llama-3-8B-Instruct" Hugging Face, <https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>
27. Yu, Hao, et al. "Evaluation of Retrieval-Augmented Generation: A Survey." arXiv preprint arXiv:2405.07437 (2024).
28. Lin, Chin-Yew. "Rouge: A package for automatic evaluation of summaries." Text summarization branches out. 2004.
29. Zhang, Tianyi, et al. "Bertscore: Evaluating text generation with bert." arXiv preprint arXiv:1904.09675 (2019).
30. Chauhan, Pratyush, et al. "Evaluating Top-k RAG-based approach for Game Review Generation." 2024 IEEE International Conference on Computing, Power and Communication Technologies (IC2PCT). Vol. 5. IEEE, 2024.