

# An Automated LLM-based Pipeline for Asset-Level Database Creation to Assess Deforestation Impact

Avanija Menon, Ovidiu Serban  
Imperial College London

## Abstract

The European Union Deforestation Regulation (EUDR) requires companies to prove their products do not contribute to deforestation, creating a critical demand for precise, asset-level environmental impact data. Current databases lack the necessary detail, relying heavily on broad financial metrics and manual data collection, which limits regulatory compliance and accurate environmental modeling. This study presents an automated, end-to-end data extraction pipeline that uses LLMs to create, clean, and validate structured databases, specifically targeting sectors with a high risk of deforestation. The pipeline introduces Instructional, Role-Based, Zero-Shot Chain-of-Thought (IRZ-CoT) prompting to enhance data extraction accuracy and a Retrieval-Augmented Validation (RAV) process that integrates real-time web searches for improved data reliability. Applied to SEC EDGAR filings in the Mining, Oil & Gas, and Utilities sectors, the pipeline demonstrates significant improvements over traditional zero-shot prompting approaches, particularly in extraction accuracy and validation coverage. This work advances NLP-driven automation for regulatory compliance, CSR (Corporate Social Responsibility), and ESG, with broad sectoral applicability.

## 1 Introduction

The European Union Deforestation Regulation (EUDR), effective December 30, 2025, mandates companies to verify that their products do not originate from recently deforested land (European Commission, 2023). With deforestation contributing 15% to global CO<sub>2</sub> emissions (ETC, 2024), industries with high environmental risks require precise asset-level tracking. However, significant data gaps persist: 30% of Forest 500 companies lack public deforestation commitments, and 85% of financial institutions lack comprehensive deforestation policies (Forest 500, 2024). Creating a physical asset database is labour intensive (CGFI, 2024), costly,

and inefficient, making regulatory compliance difficult and limiting researchers’ ability to develop accurate environmental impact models.

Due to their substantial contributions to environmental degradation, we focus on three high-risk sectors—Mining, Oil & Gas, and Utilities. Mining drives deforestation through surface extraction and infrastructure expansion, often leading to forest loss within a 50 km radius (Bradley, 2020). Oil & Gas exploration accelerates deforestation, particularly in biodiversity hotspots like the Amazon, where oil extraction disrupts ecosystems (Finer et al., 2008; Amazon Watch, 2016). Utilities, especially hydroelectric projects, contribute to deforestation (IntegrityNext, 2024) through extensive land clearing for dams and power infrastructure, with continued expansion affecting forested areas despite the shift to renewable energy (Rosenberg et al., 2000; Imperiale et al., 2023).

Our research makes several key contributions: (1) We develop a novel LLM-based pipeline (Figure 1) that transforms unstructured SEC EDGAR filings into structured datasets, improving transparency in environmental monitoring. (2) We introduce Instructional, Role-Based, Zero-Shot Chain-of-Thought (IRZ-CoT) prompting, a technique that enhances the accuracy of entity extraction, particularly for complex asset-related information. (3) We conduct a comparative analysis of LLMs and a traditional Named Entity Recognition (NER) model, evaluating their effectiveness in domain-specific data extraction. (4) To ensure data integrity, we implement a three-step database cleaning process, which includes foundational standardisation, asset similarity consolidation using statistical methods, and LLM-assisted refinement. (5) We propose Retrieval-Augmented Validation (RAV), which integrates real-time web data to enhance dataset reliability and address gaps in existing databases. (6) Finally, the resulting datasets are visualised through company-specific dashboards, providing detailed

insights into each company’s database.

This work advances NLP-driven environmental data automation, providing a scalable framework for regulatory compliance, sustainability analysis, and asset-based deforestation tracking.

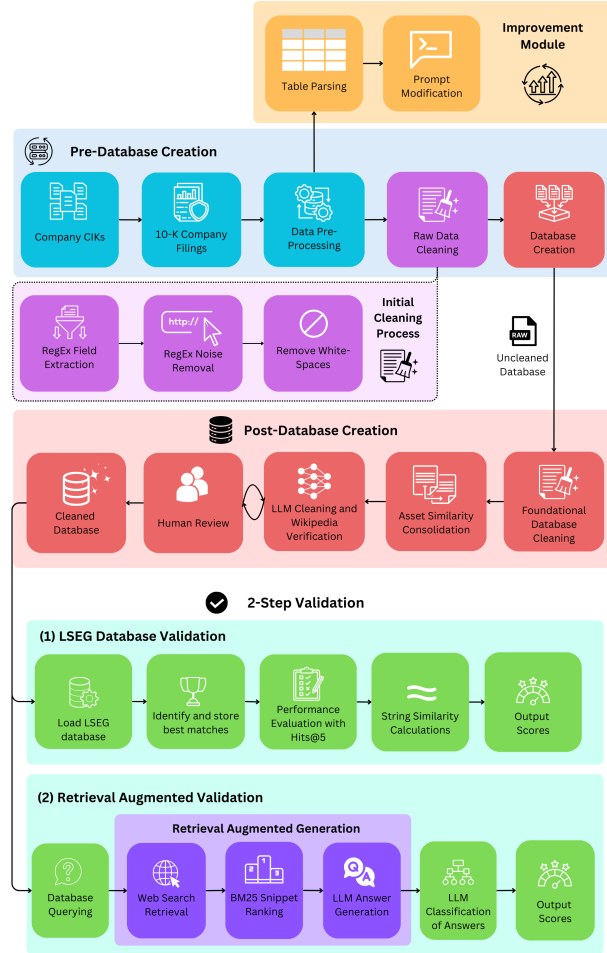


Figure 1: System design of end-to-end LLM-based pipeline designed to handle systematic data extraction, structured database creation, cleaning and validation, and the improvement module to increase validation coverage.

## 2 Background and Related Work

Existing asset-level databases rely on satellite imagery, geospatial data, and web sources with manual collection and validation (CGFI, 2024), but manual validation limits scalability. Early rule-based and template-based entity extraction methods lacked adaptability. At the same time, statistical models like Markov Logic Networks (MLNs) and Matrix Factorization improved relation extraction but were computationally expensive (Abdurehim et al., 2020; Cergani and Miettinen, 2013). Machine learning (ML) approaches, including bootstrapping (Zhang, 2009) and clustering (Tuo and

Yang, 2023), automated pattern recognition but required large labeled datasets. Deep learning models such as CNNs, RNNs, and LSTMs enhanced feature extraction and are the dominant approach (Tuo and Yang, 2023). Hybrid models, combining deep learning and traditional ML (e.g., BiLSTMs+ CNNs), to capture long-distance dependencies were explored by Zheng et al (Zheng et al., 2017).

LLMs revolutionised entity and relation extraction, enabling zero-shot and few-shot learning. Structured prompting techniques, such as Pipeline Chain-of-Thought (Pipeline-COT), enhance accuracy by breaking tasks into reasoning steps (Zhao et al., 2023). ML and NLP techniques have been widely applied in healthcare, finance, and legal domains. Transformer-based models like LegalBERT (Chalkidis et al., 2020), BioBERT (Lee et al., 2019), and SciBERT (Beltagy et al., 2019) improve clinical text analysis and regulatory compliance. However, fine-tuning remains computationally expensive, making zero-shot LLM approaches more practical. GPT-based models like GPT-NER incorporate self-verification to reduce hallucinations (Wang et al., 2023), while ChatGPT and REBEL enable structured knowledge extraction (Trajanoska et al., 2023). This study builds on these advancements, introducing Instructional, Role-Based, Zero-Shot Chain-of-Thought (IRZ-CoT) prompting to enhance structured data extraction from SEC EDGAR filings.

Traditional SEC EDGAR processing relies on RegEx-based tools like LexNLP, which efficiently parse filings (Bommarito et al., 2018). Prior keyword extraction and manual annotation work, such as the KPI-EDGAR dataset, remains labour-intensive and challenging to scale (Deußer et al., 2022).

Despite NLP advancements, limited work has been done on developing a fully automated pipeline that integrates data extraction, database creation, cleaning, and validation. This research bridges that gap by implementing an LLM-driven end-to-end pipeline, introducing Retrieval-Augmented Validation (RAV) to improve accuracy and robustness. By combining LLM-assisted extraction, structured prompts, and multi-step validation, this study delivers a scalable asset-tracking and environmental impact analysis solution, advancing AI-driven automation for regulatory data processing.

### 3 Data Acquisition and Processing

#### 3.1 Data Source

This study uses publicly available SEC EDGAR 10-K filings from fifteen Mining, Oil & Gas, and Utilities companies. These legally mandated reports provide standardised, reliable, and accurate data on company operations, finances, and environmental impact. Unlike 10-Q and 8-K reports, which offer limited asset details, 10-K filings comprehensively cover physical assets, expenditures, and disclosures. News and social media data were excluded due to bias, noise, and lack of granularity. SEC filings ensure factual accuracy, regulatory compliance, and ethical data sourcing, minimising legal and privacy concerns.

#### 3.2 Data Extraction

We collected 10-K filings from 2022 to 2024 using the secEDGAR Python library (Moody et al., 2024), which allows efficient bulk downloads based on company stock tickers and Central Index Keys (CIKs). This method streamlines data acquisition, eliminating the need for custom web scraping scripts while ensuring robust datasets across the selected sectors. The companies focused on are given in Table 3 in Appendix A.1.

The pre-processing workflow extracts metadata (company names, filing dates, form types, and content), cleans text using BeautifulSoup to remove HTML tags and irrelevant elements, and structures data into SQLite databases per company. This ensures efficient management, querying, and retention of meaningful content for analysis.

### 4 Database Creation

#### 4.1 Chunk-based Querying Technique

We adopt a chunk-based querying technique to manage the extensive length of SEC EDGAR filings. This method involves splitting documents into 1024-token chunks with a 20-token overlap to maintain contextual continuity. Sentence-level splitting ensures semantic coherence, preventing the disruption of key information. Chunking optimises memory usage, enables parallel processing, and enhances entity recognition by allowing LLMs to focus on specific, contextually rich segments. This approach also facilitates error identification and correction, improving the efficiency and scalability of the data processing pipeline.

#### 4.2 Comparison of LLM and NER Outputs

We compare the performance of 4-bit quantised Ollama instruct models, specifically Mistral-7B, Llama 3, and Gemma 2, against a traditional Named Entity Recognition (NER) model: dslim/bert-large-NER (Devlin et al., 2018; Tjong Kim Sang and De Meulder, 2003). Instruct models, fine-tuned for instruction-based tasks, demonstrate superior contextual understanding and precise entity extraction (Chung et al., 2022; Hu et al., 2024), which could be used for structured documents like SEC filings. The use of 4-bit quantisation significantly reduces memory and computational requirements while maintaining performance, enabling efficient large-scale deployment without extensive hardware upgrades (Banner et al., 2019; Dettmers et al., 2023). These models minimise irrelevant responses, ensuring more accurate asset identification. We convert the text data into embeddings using the SentenceTransformer model, specifically the paraphrase-MiniLM-L6-v2 variant (Reimers and Gurevych, 2019). Gemma 2 consistently outperforms the NER model on cosine similarity metrics, achieving higher precision and recall, with the highest cosine similarity for both locations (0.7702) and organisations (0.7461), indicating strong alignment with ground truth data.

Error analysis reveals that LLMs are more effective in capturing nuanced entity relationships, while the NER model often fragments entities or misses domain-specific terms. Detailed performance metrics are provided in Table 4 in Appendix A.3, where Gemma 2 outperforms both Mistral-7B and Llama 3.

As shown in Table 5 in Appendix A.4, qualitative error analysis highlights common issues such as fragmented entity recognition in the NER model and occasional hallucination in LLM outputs. While Mistral-7B and Llama 3 struggled with consistency, Gemma 2 demonstrated more reliable extraction, particularly in complex texts.

#### 4.3 Ground Truth Creation

We manually curated a ground truth dataset from 30 chunks of Alcoa Corporation’s 2022 filings to evaluate extraction accuracy. This dataset includes detailed annotations of physical assets, their locations, ownership structures, and associated commodities. Manual annotation ensures high accuracy, providing a robust benchmark for model evaluation. While slightly labour-intensive, this process estab-

lishes a reliable foundation for assessing model performance. In future work, we recommend exploring automated ground truth generation using advanced models like GPT-4, which could enhance scalability and reduce annotation costs.

#### 4.4 LLM Selection

We assessed multiple LLMs using evaluation metrics such as cosine and jaccard similarities, precision, recall, and F1 score. Gemma 2 emerged as the top performer, excelling in quantitative and qualitative analyses. Its superior performance is attributed to its ability to maintain semantic coherence and accurately extract domain-specific entities. As presented in Table 1, Gemma 2 achieved the highest scores across all evaluation metrics. This performance consistency and efficient resource utilisation led to its selection for further experimentation within the data pipeline.

Model	Similarity Metrics		Evaluation Metrics		
	Cosine	Jaccard	Precision	Recall	F1 Score
Mistral-7B	0.64	0.41	0.54	0.58	0.60
Llama 3	0.64	0.43	0.58	0.64	0.59
Gemma 2	0.68	0.44	0.63	0.62	0.60

Table 1: Performance comparison of Mistral-7B, Llama 3, and Gemma 2 across five metrics.

#### 4.5 Prompt Engineering

Prompt engineering plays a crucial role in optimising data extraction. We developed the Instructional, Role-Based, Zero-Shot Chain-of-Thought (IRZ-CoT) prompting technique through iterative refinement. This method improves extraction accuracy by providing LLMs with domain-specific instructions, structured reasoning steps, and role-based guidance. IRZ-CoT reduces hallucination, enhances the extraction of complex attributes, and minimises the need for extensive post-processing.

Common issues encountered with different prompting techniques revealed key challenges, such as hallucination in one-shot and few-shot methods, incorrect classification in zero-shot, and verbosity in generated knowledge prompting. Specifically, zero-shot prompting often led to the misclassification of financial terms as physical assets, while few-shot techniques introduced hallucinated entities. Role-based and instructional prompting significantly improved specificity and reduced errors, but IRZ-CoT demonstrated the best balance between accuracy and efficiency.

Performance metrics for prompt engineering

techniques, illustrated in Figure 2, show that IRZ-CoT achieved the highest scores in precision and recall. Additionally, Figure 6 in Appendix A.6 highlights IRZ-CoT’s computational efficiency, requiring significantly less processing time than methods like generated knowledge prompting.




Figure 2: Comparison of different prompt engineering techniques across various evaluation metrics

#### 4.6 Experimental Evaluation of LLM Ensemble Methods

We evaluated three LLM ensemble methods to enhance robustness: Ensemble Averaging with Majority Voting (EAMV), Weighted Majority Voting Ensemble (WMVE), and Stacking Ensemble with Meta-Learning (SEML). EAMV improves stability by aggregating predictions from multiple LLMs and selecting the most common output, reducing variance. WMVE assigns higher weights to models with superior performance, prioritising predictions from more accurate models, particularly favouring Gemma 2. SEML utilises a meta-learner—logistic regression—to combine outputs from different LLMs, optimising predictive accuracy and achieving the highest F1-score (Figure 7 in Appendix A.7). However, SEML significantly increased processing time—nearly 20-fold compared to single-model approaches (Figure 8 in Appendix A.8). Due to computational constraints, we selected the more efficient IRZ-CoT approach with Gemma 2 as the primary model for the final pipeline.

### 5 Database Cleaning

#### 5.1 Foundational Data Cleaning and Standardization

The first phase focuses on refining raw data to establish a solid foundation for further processing. We use regular expression (Regex) patterns to extract key entity data, including asset types, locations, ownership details, and commodities. This automated approach ensures consistent data extraction



from large volumes of text. Post-extraction, we remove extraneous characters, such as surplus quotes and brackets, to prevent data distortion.

Duplicates are identified and consolidated, with corresponding information merged into single records. For example, multiple entries for an oil well in different locations are grouped, reducing redundancy. Ownership data is standardised by normalising company names (e.g., consolidating "NEM," "Newmont," and "Newmont Corporation"). At the same time, geographic terms are unified (e.g., "United States of America," "US," and "U.S.A." standardised to "USA"). We also refine the 'location' column, extracting country names into a new 'Countries' column to support consistent geographic analysis. Finally, rows with empty 'physical asset' entries are removed to maintain database relevance.

## 5.2 Asset Similarity Consolidation

After initial cleaning, we address semantic similarities among physical asset entries. To consolidate such similarities, we use TF-IDF (Term Frequency-Inverse Document Frequency) vectorisation and cosine similarity.

TF-IDF quantifies the relevance of words within documents, and cosine similarity identifies semantic overlap. We set a similarity threshold of 0.5; entries meeting or exceeding this threshold are grouped and merged, preserving unique information while eliminating redundancy. This method is computationally efficient and effective for identifying similar assets, although it has limitations, such as sensitivity to synonyms. Despite these, TF-IDF and cosine similarity offer a pragmatic balance of accuracy and efficiency for large-scale datasets.

## 5.3 LLM-Assisted Database Cleaning

The final cleaning phase leverages the capabilities of Gemma 2 to address issues beyond the reach of traditional methods, as the previous steps still outputted unnormalised information amongst other issues. Using a domain-specific prompt, the LLM performs tasks such as converting chemical symbols (e.g., "Au" to "Gold"), standardising text, eliminating redundant punctuation, and verifying locations against Wikipedia.

This iterative process involves LLM-driven cleaning followed by human review. Any inconsistencies trigger prompt adjustments, enhancing the LLM's performance in subsequent iterations. The LLM also identifies countries from location data

when not explicitly stated, verified through cross-referencing with Wikipedia to ensure accuracy.

By automating complex tasks and reducing manual effort, LLM-assisted cleaning improves data quality, consistency, and scalability, making it an effective strategy for managing large datasets.

# 6 Database Validation

## 6.1 Validation with LSEG Databases

We validated our databases against established LSEG Workspace databases ([London Stock Exchange Group \(LSEG\), 2024](#)), focusing on the 'Mines', 'Oil Refineries', and 'Power Generation' datasets. This validation process involved data preprocessing, where we standardised text to lowercase and filtered irrelevant entries, such as excluding closed or abandoned assets. This step ensured uniformity and minimised discrepancies related to case sensitivity. Subsequently, we used the rapidfuzz library to find similar entries between our database entries and LSEG data. A similarity threshold of 0.6 was applied to identify potential matches, which helped us find the best match from the list of matches. We then used the Hits@5 metric to determine how frequently correct matches appeared within the top five candidates for each attribute (physical asset, ownership, commodity, and country). The Hits@5 score measures the consistency of our matching algorithm by averaging successful matches across all entities, assessing performance beyond the top result. Identified matches are then validated using five more metrics (Partial Match Score (Partial Ratio), Jaccard Similarity, Cosine Similarity, Dice-Sørensen Similarity Coefficient and Normalised Levenshtein Distance), comparing entity similarities with the LSEG database. Detailed similarity scores across physical asset name, ownership, commodity, and country are averaged into an overall attribute similarity score, quantifying dataset alignment. This validation ensures the reliability of our matching algorithm.

## 6.2 Retrieval-Augmented Validation (RAV)

LSEG Workspace databases lack comprehensive data for complete validation, necessitating an additional verification layer to ensure completeness and accuracy. To address these gaps, we developed Retrieval-Augmented Validation (RAV).

RAV integrates real-time web search capabilities using the Google Custom Search Engine (CSE)

API (Google Developers, 2024) to retrieve current information on physical assets. The retrieved snippets are ranked using the BM25 algorithm, which prioritises documents based on relevance, incorporating term frequency and document length normalisation. This ensures that the most pertinent information is considered for validation purposes.

RAV uses a dual-LLM framework where Llama 3 generates web-based answers, and Gemma 2 is tasked with classifying these answers strictly against the database entries. This separation mitigates potential biases arising from using a single model for generation and evaluation. Llama 3 efficiently retrieves concise, relevant information from web sources, while Gemma 2 assesses the similarity between this information and the existing database entries. The LLM-assisted validation relies on a binary classification approach where Gemma 2 outputs a ‘yes’ if the web-derived and database information are similar and a ‘no’ otherwise. This stringent evaluation ensures high reliability, reducing the risk of false positives in validation.

Contrary to complex instructional prompts used in earlier phases, we discovered that simple prompts significantly improved LLM classification accuracy. Initially, using detailed prompts resulted in low similarity scores, averaging around 0.15, with frequent misclassifications. After simplifying the prompts to a single-line instruction asking the LLM to classify answers as similar or dissimilar (see Appendix A.13), we observed a substantial improvement, with scores increasing by approximately 0.28. This reduction in cognitive load enhanced the model’s ability to determine similarities accurately. However, some misclassifications remain, mainly when subtle semantic differences exist between the database entries and web-sourced information.

RAV automates asset validation by integrating web data with traditional databases, enhancing reliability for downstream analysis. While advanced RAG methods like FLARE (Jiang et al., 2023) offer sophisticated retrieval, their complexity and resource demands outweigh the benefits for this project. Our BM25-based RAV remains practical and effective, with potential for future refinement.

## 7 Results

### 7.1 LSEG Database Validation Results

We validated our databases against LSEG Workspace datasets, including ‘Mines,’ ‘Oil Refineries,’ and ‘Power Generation,’ using six similarity metrics: Partial Match Score, Jaccard Similarity, Cosine Similarity, Dice-Sørensen Similarity Coefficient, Normalised Levenshtein Distance, and Hits@5. These metrics assessed alignment across physical assets, ownership, commodities, and country data. Figure 3 shows the partial match scores. The full results are shown in Figure 9 in Appendix A.9.




Figure 3: Partial match scores from the LSEG database validation. The dotted lines separate the sectors, where the sectors are mining, oil & gas, and utilities, respectively.

The mining sector demonstrated strong alignment, with high partial match scores for companies such as AA (0.95), FCX (0.88), and NEM (0.83), reflecting data consistency. However, the oil & gas sector showed mixed results, where CVX (0.72) and MPC (0.74) achieved moderate alignment, but XOM exhibited lower scores (Jaccard Similarity: 0.24), suggesting inconsistencies in asset classification. Utilities displayed moderate alignment, with EXC and D achieving partial match scores of 0.93 and 0.79, respectively. Ownership data varied significantly, with CVX achieving near-perfect alignment (1.00), while XOM had lower similarity (Jaccard Similarity: 0.43), likely due to differences in how joint ventures and subsidiaries were recorded. Commodity data showed the most significant discrepancies, with many companies, such as AA, registering low Jaccard and Cosine Similarities (0.00), possibly due to differences in classifying primary and secondary commodities. In contrast, country data was generally consistent, with companies like FCX achieving perfect alignment (Partial

Match Score: 1.00), though some discrepancies were observed in SCCO (Jaccard Similarity: 0.67).

The error analysis revealed key challenges. Ownership discrepancies arose due to variations in recording structures, where our databases captured joint ventures while LSEG focused on primary controlling entities. Standardising ownership classification could improve future alignment. Commodity misalignments resulted from differences in listing primary versus secondary commodities, suggesting a need to refine entity extraction prompts and separate commodity categories. Minor inconsistencies in country data, such as listing "USA" versus "California," highlight the importance of hierarchical structuring with separate fields for city, region, and country to enhance accuracy.

## 7.2 Coverage Calculation

We evaluate database coverage by measuring the proportion of physical assets and their attributes in our constructed database that match those in LSEG. This assessment ensures comprehensiveness, usability, and accuracy while identifying areas for improvement in our extraction pipeline. Coverage is computed as  $\text{Coverage Score} = \left( \frac{N_m}{N_L} \right) \times 100$ , where  $N_m$  represents the number of matched physical assets between our database and LSEG, and  $N_L$  is the total number of physical assets in LSEG. The computed coverage scores, shown in Table 7 in Appendix A.10, indicate that the mining sector has better coverage than oil & gas and utilities.

Manual inspection of SEC EDGAR filings reveals that lower coverage in oil & gas and utilities stems from improper table parsing, as many assets are listed in tabular formats rather than continuous text. To address this, we integrated a table parsing module using LlamaIndex, which processes HTML tables as structured data instead of narrative text. This significantly improved extraction accuracy, particularly in oil & gas, where assets were previously missed. Figure 10 in Appendix A.11 demonstrates this enhancement.

## 7.3 Retrieval-Augmented Validation (RAV) Results

Table 8 in Appendix A.12 presents the results of Retrieval-Augmented Validation (RAV), comparing database responses with real-time web data. Similarity scores range from 0.31 to 0.57, indicating moderate alignment, with the Oil & Gas sector performing slightly better due to more transparent

regulatory disclosures. Notably, OXY is absent from Table 8 since it only contains unnamed assets, which RAV cannot validate.

Mining sector scores hover around 0.4, suggesting uniform discrepancies, likely due to outdated or incomplete records. The Oil & Gas sector shows slightly higher alignment, with companies like MPC and COP exceeding 0.5, possibly due to stringent regulatory reporting. However, frequent asset transfers contribute to inconsistencies. The Utilities sector exhibits the widest score range, from 0.31 (EXC) to 0.57 (NEE), reflecting differences in data transparency. NEE's higher score suggests more consistent asset records, likely due to better data management.

### 7.3.1 Error Analysis

Ownership mismatches arose from differing data granularity. Our database captured joint ventures and minority stakeholders, whereas web sources listed only primary entities, leading to unfair scores of 0. A weighted scoring system could better account for partial matches.

Location mismatches often resulted from implicit references in web snippets. For instance, the Bath County Power Station was correctly labeled as USA in our database, but the web snippet lacked an explicit country mention, receiving a score of 0. Similarly, Chino Mine was recorded as USA, while web sources specified New Mexico, USA. A hierarchical scoring approach would improve accuracy by recognising different levels of geographic detail.

Commodity discrepancies occurred because web data often listed only primary commodities, while our database included by-products. For example, Grasberg Mine was recorded as producing copper, gold, silver, and molybdenum, whereas web results mentioned only silver. Categorising commodities into primary and secondary groups through prompt refinement would help resolve this.

## 7.4 Total Validation Coverage

To assess RAV's impact, we compute the total validation coverage, which measures the proportion of assets validated through both LSEG database validation and RAV. Total validation coverage is computed as  $\left( \frac{N_v}{N_t} \right) \times 100$ , where  $N_v$  represents the number of assets validated, and  $N_t$  is the total number of assets in the constructed database.

Table 2 presents validation coverage for each company, comparing LSEG-only validation to com-

bined LSEG and RAV validation. Occidental Petroleum (OXY) is excluded due to the absence of company-specific information in the LSEG database and the constructed dataset containing only general assets (e.g., natural gas fields).

Since our validation applies only to named assets, general assets remain largely unverified. While extrapolating validation to unnamed assets could extend coverage, this introduces risks to accuracy and completeness. Notably, RAV significantly increases coverage, underscoring its role in enhancing database robustness by validating assets absent from LSEG.

Coverage varies across companies; D achieves the highest at 33.33%, while COP has the lowest at 6.43%, reflecting differences in named asset proportions. Lower coverage suggests a higher proportion of unnamed assets, highlighting gaps in the current validation process.

Sector	Company	Validation Coverage	
		LSEG Database	LSEG Database + RAV
Mining	AA	6.96%	20.00%
	SCCO	7.27%	7.27%
	FCX	7.35%	21.32%
	HL	17.11%	21.05%
	NEM	19.51%	30.89%
Oil and Gas	CVX	7.61%	20.65%
	XOM	7.69%	23.08%
	MPC	15.12%	25.58%
	COP	0.71%	6.43%
Utilities	D	6.06%	33.33%
	ED	1.94%	17.48%
	DUK	3.90%	23.38%
	EXC	9.57%	24.35%
	NEE	0%	10.77%

Table 2: Validation coverage comparison using LSEG databases alone versus LSEG databases with RAV.

As regulatory demands like the EUDR grow, the need for automated, comprehensive databases will increase. Our LLM-based pipeline can adapt to these demands, improving ESG and CSR compliance. The feedback loop (Figure 4) from regulatory success will drive continuous improvements in data quality and database creation techniques, shaping the future of environmental data management.

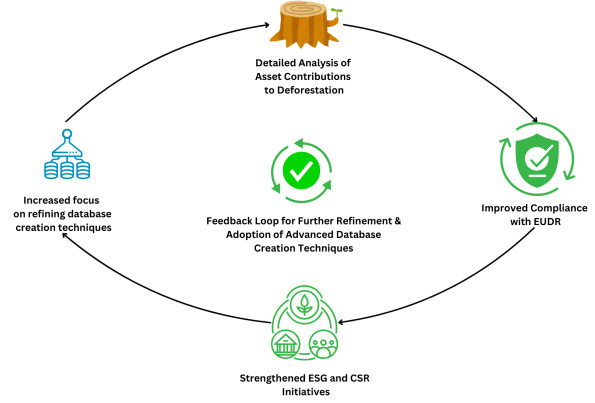


Figure 4: A feedback loop linking physical asset database creation with improved compliance and ESG initiatives, driving continuous refinement.

## 8 Conclusion

This study developed an end-to-end LLM-based pipeline for extracting, structuring, cleaning, and validating physical asset data from SEC EDGAR filings, providing a scalable and automated approach to asset tracking and environmental impact analysis. The pipeline follows a structured architecture, incorporating data acquisition and processing, entity extraction, database creation, database cleaning and a two-step validation framework: LSEG database validation followed by Retrieval-Augmented Validation (RAV). The table parsing improvement module also significantly increased validation coverage, enhancing database completeness. Using LLMs for automation, this study advances structured data extraction and validation, laying the foundation for more efficient regulatory compliance and environmental data management. As regulatory demands evolve, LLM-based techniques will play an increasingly critical role in ensuring accurate, structured asset tracking.

## 9 Limitations

Despite its success, this study has several limitations. The dataset was limited to fifteen companies across three sectors, using only three years of SEC filings, which may restrict the generalisability of the findings. Additionally, the project relied heavily on SEC 10-K filings, excluding potential insights from 8-K, 10-Q, and other reports that may contain relevant asset data. While the IRZ-CoT prompt engineering technique significantly improved entity extraction, its effectiveness in other domains or regulatory environments remains untested.

The integration of LLMs introduces maintenance



and adaptation challenges, especially as newer models may require retuning for continued effectiveness. Biases in pre-trained data could also impact extraction accuracy, particularly in financial and environmental sectors. The RAV validation process depends on web search quality, meaning incomplete or inaccurate online data could affect validation reliability. Although automation reduced manual effort, human verification was necessary for ground truth dataset creation and qualitative assessments. Additionally, the pipeline primarily focuses on historical data, making real-time asset tracking and monitoring challenging.

To address these limitations, future improvements should refine prompt engineering to extract more named assets, reducing reliance on proxies. Expanding the dataset to additional industries and historical filings will enhance coverage and validation effectiveness. Fine-tuning LLMs using domain-specific databases could enhance extraction accuracy. Additionally, dynamic prompt engineering combined with reinforcement learning — where the model is rewarded for accurate extractions and penalized for errors — could help the system adapt more effectively to different types of company disclosures.

RAV significantly increases validation coverage. It could further benefit from integrating advanced web search APIs and adopting a weighted scoring system to account for partial matches, improving validation granularity. Broader data issues also affected validation when web search results were incomplete or irrelevant, leading to default scores of 0. Expanding search queries, integrating multiple search engines, and refining data extraction techniques could improve RAV’s robustness. Lastly, enhancing automated table parsing would improve structured data extraction from SEC filings, particularly in appendices and financial disclosures where asset details are often tabulated. This could be achieved by automated prompt tailoring, where an LLM (e.g., GPT-4) identifies asset examples from parsed tables and incorporates them into IRZ-CoT prompts.

## 10 Ethics Statement

This project adheres to the ACM Code of Ethics (ACM, 2018) and the Ethics Guidelines for Trustworthy AI (European Commission, 2019), emphasizing transparency, fairness, accountability, and technical robustness. It relies on publicly available

SEC EDGAR filings, ensuring legal compliance and data integrity. Ethical data acquisition is upheld by adhering to data source terms and conditions, preventing legal conflicts. The project contributes to environmental sustainability by assessing industrial deforestation impacts, supporting global climate initiatives, and reinforcing corporate social responsibility (CSR) through ESG compliance. Accuracy and fairness are prioritized via a two-step validation pipeline, minimizing misrepresentations and addressing LLM biases through transparency and human oversight. By maintaining industry best practices in data processing, the project ensures reliability, robustness, and responsible AI development for ethical and accurate asset tracking.

## Acknowledgments

We would like to express our sincere gratitude to the LSEG (London Stock Exchange Group) team for providing a free license to the LSEG Workspace and for their invaluable guidance throughout the course of this research. Their support significantly contributed to the quality and depth of this work.

## References

- Peride Abdurehim, Turdy Tohti, and Askar Hamdulla. 2020. [A short review of relation extraction methods](#). In *2020 13th International Conference on Intelligent Computation Technology and Automation (ICICTA)*, pages 18–22.
- ACM. 2018. [Acm code of ethics and professional conduct](#).
- Amazon Watch. 2016. [The amazon crude: How american consumers fuel deforestation and human rights abuses in south america](#). Technical report, Amazon Watch, Oakland, CA.
- Ron Banner, Yury Nahshan, Elad Hoffer, and Daniel Soudry. 2019. [Post-training 4-bit quantization of convolution networks for rapid-deployment](#).
- Iz Beltagy, Arman Cohan, and Kyle Lo. 2019. [Scibert: Pretrained contextualized embeddings for scientific text](#). *CoRR*, abs/1903.10676.
- M. Bommarito, D. Katz, and E. M. Detterman. 2018. [LexNLP: Natural Language Processing and Information Extraction For Legal and Regulatory Texts](#). *InfoSciRN: Legal Informatics (Topic)*.
- Sian Bradley. 2020. [Mining’s impacts on forests: Aligning policy and finance for climate and biodiversity goals](#). Technical report, Chatham House, London, UK.

- Ervina Cergani and Pauli Miettinen. 2013. [Discovering relations using matrix factorization methods](#). In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13*, page 1549–1552, New York, NY, USA. Association for Computing Machinery.
- CGFI. 2024. [Geoasset project: Petrochemicals](#).
- Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, and Ion Androutsopoulos. 2020. [LEGAL-BERT: the muppets straight out of law school](#). *CoRR*, abs/2010.02559.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#).
- Tobias Deußer, Syed Musharraf Ali, Lars Hillebrand, Desiana Nurchalifah, Basil Jacob, Christian Bauckhage, and Rafet Sifa. 2022. [Kpi-edgar: A novel dataset and accompanying metric for relation extraction from financial documents](#). In *2022 21st IEEE International Conference on Machine Learning and Applications (ICMLA)*. IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- ETC. 2024. [Financing the transition: The costs of avoiding deforestation](#). Technical report, Energy Transitions Commission, London, UK.
- European Commission. 2019. [Ethics guidelines for trustworthy ai](#).
- European Commission. 2023. [Regulation on deforestation-free products](#).
- Matt Finer, Clinton N. Jenkins, Stuart L. Pimm, Brian Keane, and Carl Ross. 2008. [Oil and gas projects in the western amazon: Threats to wilderness, biodiversity, and indigenous peoples](#). *PLoS ONE*, 3(8):e2932.
- Forest 500. 2024. [2024: A decade of deforestation data](#).
- Google Developers. 2024. [Custom search json api overview](#).
- Hanxu Hu, Simon Yu, Pinzhen Chen, and Edoardo M. Ponti. 2024. [Fine-tuning large language models with sequential instructions](#).
- Francesca Imperiale, Simone Pizzi, and Stella Lippolis. 2023. [Sustainability reporting and esg performance in the utilities sector](#). *Utilities Policy*, 80:101468.
- IntegrityNext. 2024. [Energy & utilities](#).
- Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Active retrieval augmented generation](#).
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. [BioBERT: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- London Stock Exchange Group (LSEG). 2024. [Lseg workspace: Data and analytics](#).
- Moody et al. 2024. [Sec-edgar: Download all companies' periodic reports, filings, and forms from edgar database](#). GitHub repository, Apache-2.0 License.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- David M. Rosenberg, Patrick McCully, and Catherine M. Pringle. 2000. [Global-scale environmental effects of hydrological alterations: Introduction](#). *BioScience*, 50(9):746–751.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Milena Trajanoska, Riste Stojanov, and Dimitar Trajanov. 2023. [Enhancing knowledge graph construction using large language models](#).
- Meimei Tuo and Wenzhong Yang. 2023. [Review of entity relation extraction](#). *Journal of Intelligent & Fuzzy Systems*, 44(5):7391–7405.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023. [Gpt-ner: Named entity recognition via large language models](#).
- Suxiang Zhang. 2009. [Entity relation extraction to free text](#). In *2009 International Conference on Natural Language Processing and Knowledge Engineering*, pages 1–5.
- Hangtian Zhao, Hakiz Yilahun, and Askar Hamdulla. 2023. [Pipeline chain-of-thought: A prompt method for large language model relation extraction](#). In *2023 International Conference on Asian Language Processing (IALP)*, pages 31–36.

Suncong Zheng, Yuexing Hao, Dongyuan Lu, Hongyun Bao, Jiaming Xu, Hongwei Hao, and Bo Xu. 2017. Joint entity and relation extraction based on a hybrid neural network. *Neurocomputing*, 257:59–66. Machine Learning and Signal Processing for Big Multimedia Analysis.

## A Appendix

### A.1 Sectors, Companies and Stock Tickers used

Sector	Company Name	Stock Ticker
Mining	Alcoa Corporation	AA
	Hecla Mining Corporation	HL
	Newmont Corporation	NEM
	Freeport Mc-Moran	FCX
	Southern Copper Corporation	SCCO
Oil & Gas	ConocoPhillips Company	COP
	Marathon Petroleum Corporation	MPC
	Chevron Corporation	CVX
	Occidental Petroleum	OXY
	Exxon Mobil Corporation	XOM
Utilities	Dominion Energy	D
	Duke Energy Corporation	DUK
	Consolidated Edison	ED
	Exelon Corporation	EXC
	NextEra Energy	NEE

Table 3: Companies used for pipeline construction, grouped by sector.

### A.2 Chunks per document in Chunk-based Querying Technique




Figure 5: Number of chunks generated per document for each company across the three sectors.

### A.3 Cosine similarities of LLM-generated outputs against NER results

Model	Location - cosine similarity	Organisation - cosine similarity
Mistral-7B	0.6945	0.6809
Llama 3	0.5987	0.7177
Gemma 2	0.7702	0.7461

Table 4: Cosine similarities of LLM-generated outputs against NER results.

#### A.4 Error Analysis of LLM-generated outputs against NER results

Model	Type	Description (LLM vs. NER Output)
Mistral-7B	Success	<b>LLM:</b> Accurately extracted locations "South Church Street, Charlotte, North Carolina" and organisations. <b>NER:</b> Fragmented these into "Securities Exchange" and "##TON." Split up locations like 'US' into 'U' and 'S'.
	Failure	<b>LLM:</b> Sometimes failed to extract any entities. <b>NER:</b> Correctly identified "Duke Energy."
	Challenge	<b>LLM:</b> Provided a list of relevant organisations but missed key locations. <b>NER:</b> Captured more entities but included irrelevant fragments like "GENERA."
LLaMA 3	Success	<b>LLM:</b> Correctly extracted "Duke Energy" without fragmentation. <b>NER:</b> Fragmented other entities and introduced spurious ones like "Per" and "Board of Directors".
	Failure	<b>LLM:</b> Sometimes missed all entities. <b>NER:</b> Extracted locations like "Central" and organisations like "Spectra Energy."
	Challenge	<b>LLM:</b> Identified relevant text but added repeated/irrelevant words. <b>NER:</b> Output was more accurate but fragmented.
Gemma 2	Success	<b>LLM:</b> Correctly extracted locations and organisations. <b>NER:</b> Fragmented entities and returned partial results.
	Failure	<b>LLM:</b> Sometimes missed all entities. <b>NER:</b> Extracted "Duke Energy" and "Cinergy."
	Challenge	<b>LLM:</b> Missed key locations. <b>NER:</b> Output included a broader range of organisations but included irrelevant fragments.

Table 5: Examples of Successes, Failures, and Challenges in entity extraction by LLMs compared to NER.



## A.5 Problem descriptions in Iterative Prompt Refinement

Prompt Name	Problem Description	Problem Example
Zero-shot	Incorrect mention of financial assets as physical assets, lack of location specificity	'asset: Consolidated Operations - Revenues', 'location: Seven countries'
One-shot	Hallucination — includes example in answer	'ownership: Freeport-McMoran'
Few-shot	Incorrect classification of physical asset, hallucination — includes example in answer	'physical asset: Granted patents (intellectual property), Registered trademarks (intellectual property)', 'physical assets: [Grasberg mine]'
Chain-of-thought	Incorrect mention of location as physical asset; provides irrelevant answers if no information is available	'physical asset: Brazil', 'relationships: [asset: ", location: ", ownership: ", commodities: "]'
Generated knowledge	Outputs overly detailed and verbose answers	'physical assets: [substantially all assets of the Company, ...]'
Prompt chaining	Lack of specificity	'physical assets: [substantially all assets of the Company]', 'physical assets: [mining operations, properties, leases]'
Role prompting	Improvement noted but lack of specificity	'physical assets: [Mining operations, processing plants, manufacturing facilities]'
Role and instructional prompting	Improvement noted but lack of specificity	'physical assets: [facilities]'

Table 6: Problem descriptions and examples for each prompting technique.

## A.6 Computational times of prompting techniques in Iterative Prompt Refinement

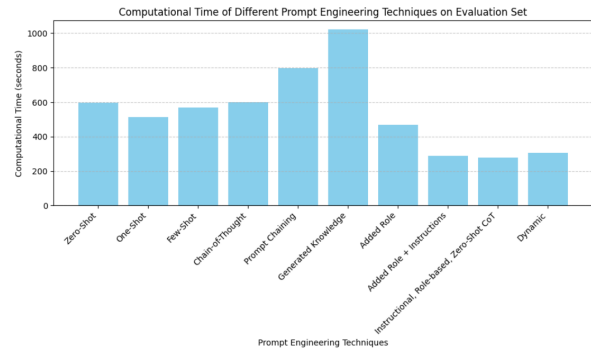


Figure 6: Computational time (in seconds) required for different prompt engineering techniques on the evaluation set.

## A.7 Results of LLM Ensemble Implementation

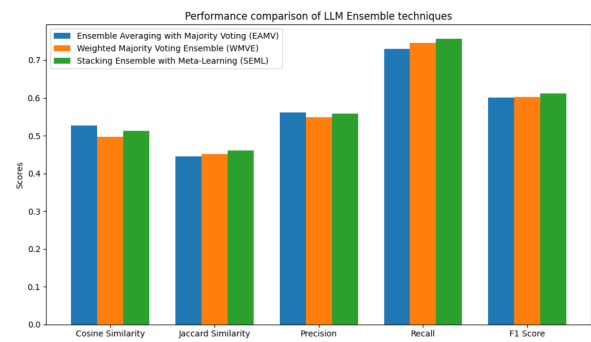


Figure 7: Performance comparison of various LLM ensemble techniques.

## A.8 Efficiency of LLM Ensemble Methods




Figure 8: Comparison of computational times with and without using LLM ensemble techniques.

## A.9 LSEG Database Validation Results

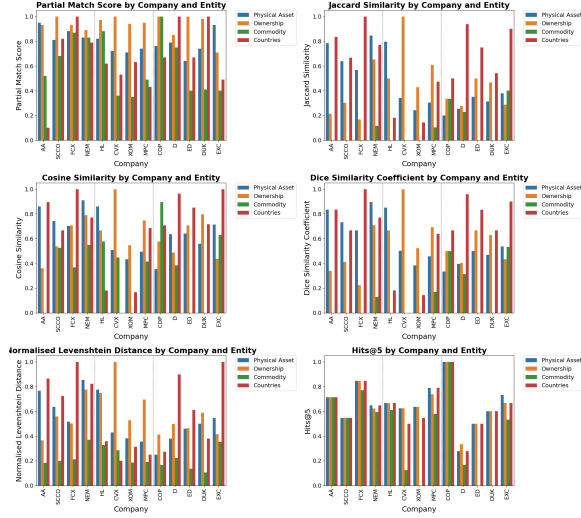


Figure 9: Results from the LSEG database validation across various metrics.

## A.10 Coverage Calculation Results

Sector	Company	Coverage
Mining	AA	71.43%
	SCCO	54.55%
	HL	66.67%
	NEM	62.16%
	FCX	92.31%
Oil & Gas	CVX	25%
	XOM	0%
	MPC	94.74%
	OXY	N/A
	COP	100%
Utilities	D	27.78%
	DUK	40%
	EXC	13.33%
	ED	50%
	NEE	N/A

Table 7: Calculated coverage percentages for databases across sectors.

## A.11 Coverage before and after implementation of table parsing

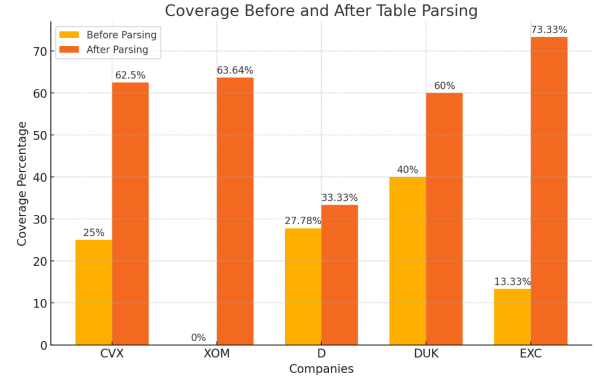


Figure 10: Coverage percentages before and after performing table parsing.

## A.12 Retrieval Augmented Validation (RAV) Results

Sector	Company	Similarity Score
Mining	AA	0.42
	SCCO	0.37
	FCX	0.43
	NEM	0.38
	HL	0.46
Oil and Gas	CVX	0.44
	XOM	0.48
	MPC	0.52
	COP	0.52
Utilities	D	0.33
	ED	0.38
	DUK	0.40
	EXC	0.31
	NEE	0.57

Table 8: Averaged classification similarity scores from Retrieval Augmented Validation (RAV).

## A.13 Prompt Library

In this section, we present the comprehensive collection of prompts utilised throughout this project for information extraction. The prompt library consists of a variety of carefully designed instructions aimed at guiding LLMs in extracting specific entities and relationships, such as physical assets, locations, ownership details, and commodities, from SEC EDGAR filings. Each prompt is tailored to enhance the performance of LLMs in different scenarios, employing techniques such as zero-shot, one-shot, few-shot prompting, and more sophisticated methods like Chain-of-Thought (CoT) reasoning and role-based prompting, before

constructing our IRZ-CoT prompt.

The prompts are categorised based on their usage in different stages of the project, including entity extraction, database cleaning and when creating an improvement module. This prompt library serves as the foundation for automating the extraction process and ensuring the reliability and accuracy of the data. Each prompt has been optimised through iterative testing and refinement to address the unique challenges posed by our use case.

#### Prompt for extracting locations and organisations when comparing LLMs and NER

```
prompt_instruction = "You are
a virtual assistant with
advanced expertise in a
broad spectrum of topics,
equipped to utilize high-
level critical thinking,
cognitive skills,
creativity, and innovation.
Your goal is to deliver
the most straightforward
and accurate answer
possible for each question,
ensuring high-quality and
useful responses for the
user."
user_prompt = f"Text: {chunk}\n
Query: Does this text
mention any locations or
organisations? If yes,
please specify them in the
following format:\n
locations: [ ]\n
organisations: [ ]"
```

#### Zero-shot

```
f"Text: {chunk}\nQuery: Does this
text mention any physical
assets, locations, ownerships,
and commodities? "
"If yes, please specify
them in the following
format:\n"
"physical assets: [ ]\n
locations: [ ]\n
ownerships: [ ]\n
commodities: [ ]\n"
"Additionally, identify
the relationships
between them,
specifying the location
of each physical asset
, their ownership
details, and
commodities. "
"Format the relationships
```

```
as follows:\n
relationships: [asset:
'', location: '',
ownership: '',
commodities: '']"
```

#### One-shot

```
f"Text: {chunk}\nQuery: Does this
text mention any physical
assets, locations, ownerships,
and commodities? "
"If yes, please specify
them in the following
format:\n"
"physical assets: [ ]\n
locations: [ ]\n
ownerships: [ ]\n
commodities: [ ]\n"
"Additionally, identify
the relationships
between them,
specifying the location
of each physical asset
, their ownership
details, and
commodities. "
"Format the relationships
as follows:\n
relationships: [asset:
'', location: '',
ownership: '',
commodities: '']"
"Here is an example:\n"
"Example:\n"
"Text: [...] Our principal
asset is the Grasberg
mine, which we
discovered in 1988.
Grasberg contains the
largest single gold
reserve and one of the
largest copper reserves
of any mine in the
world. Our principal
operating subsidiary is
PT Freeport Indonesia,
a limited liability
company organized under
the laws of the
Republic of Indonesia
and incorporated in
Delaware. [...]"
"Query: Does this text
mention any physical
assets, locations, and
ownerships?\n"
"physical assets: [
Grasberg mine]\n
locations: [Sudirman
Mountain Range, Papua,
Indonesia]\nownerships:
[Republic of Indonesia
, Delaware]\n[
commodities: copper,
gold]\n"
```

```
"relationships: [asset: '
  Grasberg mine',
  location: 'Indonesia',
  ownership: 'PT Freeport
  Indonesia',
  commodities: 'copper',
  'gold']\n\n"
```

### Few-shot

```
f"Text: {chunk}\nQuery:
  Does this text mention
  any physical assets,
  locations, ownerships,
  and commodities? "
"If yes, please specify
  them in the following
  format:\n"
"physical assets: [ ]\
nlocations: [ ]\
nownerships: [ ]\
ncommodities: [ ]\n"
"Additionally, identify
  the relationships
  between them,
  specifying the location
  of each physical asset
  , their ownership
  details, and
  commodities. "
"Format the relationships
  as follows:\
nrelationships: [asset:
  '', location: '',
  ownership: '',
  commodities: '']"
"Here are some examples:\n
"
"Example:\n"
"Text: [...] Our principal
  asset is the Grasberg
  mine, which we
  discovered in 1988.
  Grasberg contains the
  largest single gold
  reserve and one of the
  largest copper reserves
  of any mine in the
  world. Our principal
  operating subsidiary is
  PT Freeport Indonesia,
  a limited liability
  company organized under
  the laws of the
  Republic of Indonesia
  and incorporated in
  Delaware. [...]"
"Query: Does this text
  mention any physical
  assets, locations, and
  ownerships?\n"
"physical assets: [
  Grasberg mine]\
nlocations: [Sudirman
  Mountain Range, Papua,
  Indonesia]\nownerships:
```

```
[Republic of Indonesia
, Delaware]\n[
commodities: copper,
gold]\n"
"relationships: [asset: '
  Grasberg mine',
  location: 'Indonesia',
  ownership: 'PT Freeport
  Indonesia',
  commodities: 'copper',
  'gold']\n\n"
```

```
"Example 2:\n"
"Text: [...] PT Freeport
  Indonesia mines,
  processes and explores
  for ore containing
  copper, gold and silver
  . It operates in the
  remote highlands of the
  Sudirman Mountain
  Range in the province
  of Papua (formerly
  Irian Jaya), Indonesia,
  which is on the
  western half of the
  island of New Guinea.
  [...]"
"Query: Does this text
  mention any physical
  assets, locations, and
  ownerships?\n"
"physical assets: [PT
  Freeport Indonesia
  Mines]\nlocations: [
  Sudirman Mountain Range
  , Papua, Indonesia, New
  Guinea]\nownerships: [
  PT Freeport Indonesia]\
ncommodities: [copper,
  gold, silver]\n"
"relationships: \n"
"[asset: 'PT Freeport
  Indonesia Mines',
  location: 'Sudirman
  Mountain Range, Papua,
  Indonesia, New Guinea',
  ownership: 'PT
  Freeport Indonesia',
  commodities: 'copper,
  gold, silver']\n"
```

```
"Example 3:\n"
"Text: [...] The Republic
  of Indonesia consists
  of more than 17,000
  islands stretching
  3,000 miles along the
  equator from Malaysia
  to Australia and is the
  fourth most populous
  nation in the world
  with over 200 million
  people. [...]"
"Query: Does this text
  mention any physical
  assets, locations, and
  ownerships?\n"
```



```
"physical assets: []\n
  nlocations: [Republic
    of Indonesia, Malaysia,
    Australia]\nownerships
    : []\ncommodities: []\n
  \n"
"relationships: \n"
```

### Chain-of-Thought (CoT) prompting

```
f"Text: {chunk}\nQuery: Let's
  think step by step. First,
  identify any physical assets
  mentioned in the text. Next,
  determine if any locations or
  ownership details are provided
  for these physical assets. Then
  , determine if the commodities
  related to the physical assets
  are provided. Finally,
  summarize the relationships
  between each physical asset,
  its location, its ownership and
  its commodity. "
  "If yes, please specify
  them in the following
  format:\n"
  "physical assets: [ ]\n
    nlocations: [ ]\n
    nownerships: [ ]\n
    ncommodities: [ ]\n"
  "Additionally, identify
  the relationships
  between them,
  specifying the location
  of each physical asset
  , their ownership
  details, and
  commodities. "
  "Format the relationships
  as follows:\n
  nrelationships: [asset:
    '', location: '',
    ownership: '',
    commodities: '']"
```

### Generated knowledge prompting

```
knowledge_prompt = (
  "You are an expert in
  analyzing texts for
  information about
  physical assets,
  locations, ownerships,
  and commodities. "
  "Provide a brief summary
  of how to identify
  these elements in a
  text and the
  relationships between
  them."
)
prompt_step1 = (
```

```
f"{generated_knowledge}\n\n"
  "You are a virtual
  assistant with
  expertise in extracting
  specific information
  from text. "
  "A physical asset is an
  asset with a
  geographical location.\n
  \n"
  f"Text: {chunk}\nQuery:
    Identify any physical
    assets mentioned in the
    text. "
  "List them in the format:\n
    nphysical assets: [ ]"
)
prompt_step2 = (
  f"{generated_knowledge}\n\n"
  "Using the extracted
  physical assets:\n"
  f"physical assets: {
    physical_assets}\n\n"
  f"Text: {chunk}\nQuery:
    Identify any locations
    mentioned in the text
    associated with the
    physical assets. "
  "List them in the format:\n
    nlocations: [ ]"
)
prompt_step3 = (
  f"{generated_knowledge}\n\n"
  "Using the extracted
  physical assets and
  locations:\n"
  f"physical assets: {
    physical_assets}\n
    nlocations: {locations
    }\n\n"
  f"Text: {chunk}\nQuery:
    Identify any ownership
    details mentioned in
    the text associated
    with the physical
    assets. "
  "List them in the format:\n
    nownerships: [ ]"
)
prompt_step4 = (
  f"{generated_knowledge}\n\n"
  "Using the extracted
  physical assets,
  locations, and
  ownerships:\n"
  f"physical assets: {
    physical_assets}\n
    nlocations: {locations
    }\nownerships: {
    ownerships}\n\n"
  f"Text: {chunk}\nQuery:
    Identify any
    commodities mentioned
    in the text associated
```

```

        with the physical
        assets. "
    "List them in the format:\
    ncommodities: [ ]"
)
prompt_step5 = (
    f"{generated_knowledge}\n\
    n"
    "Using the extracted
    physical assets,
    locations, ownerships,
    and commodities:\n"
    f"physical assets: {
    physical_assets}\
    nlocations: {locations
    }\nownerships: {
    ownerships}\
    ncommodities: {
    commodities}\n\n"
    "Text: {chunk}\nQuery:
    Identify the
    relationships between
    the physical assets,
    locations, ownerships,
    and commodities. "
    "Format the relationships
    as follows:\
    nrelationships: [asset:
    '', location: '',
    ownership: '',
    commodities: '']"
)

```

### Prompt Chaining

```

prompt_step1 = (
    f"Text: {chunk}\nQuery:
    Does this text mention
    any physical assets? "
    "If yes, please specify
    them in the following
    format:\n"
    "physical assets: [ ]"
)

prompt_step2 = (
    f"physical assets: {
    physical_assets}\n"
    f"Text: {chunk}\nQuery:
    Does this text mention
    any locations
    associated with the
    physical assets? "
    "If yes, please specify
    them in the following
    format:\n"
    "locations: [ ]"
)

prompt_step3 = (
    f"physical assets: {
    physical_assets}\
    nlocations: {locations
    }\n"
    f"Text: {chunk}\nQuery:

```

```

    Does this text mention
    any ownership details
    associated with the
    physical assets? "
    "If yes, please specify
    them in the following
    format:\n"
    "ownerships: [ ]"
)

prompt_step4 = (
    f"physical assets: {
    physical_assets}\
    nlocations: {locations
    }\nownerships: {
    ownerships}\n"
    f"Text: {chunk}\nQuery:
    Does this text mention
    any commodities
    associated with the
    physical assets? "
    "If yes, please specify
    them in the following
    format:\n"
    "commodities: [ ]"
)

prompt_step5 = (
    f"physical assets: {
    physical_assets}\
    nlocations: {locations
    }\nownerships: {
    ownerships}\
    ncommodities: {
    commodities}\n"
    f"Text: {chunk}\nQuery:
    Identify the
    relationships between
    the physical assets,
    locations, ownerships,
    and commodities. "
    "Format the relationships
    as follows:\
    nrelationships: [asset:
    '', location: '',
    ownership: '',
    commodities: '']"
)

```

### Role prompting

```

prompt_instruction = (
    "You are a virtual
    assistant with advanced
    expertise in a broad
    spectrum of topics,
    equipped to utilize
    high-level critical
    thinking, cognitive
    skills, creativity, and
    innovation.\n"
    "Your goal is to deliver
    the most
    straightforward and
    accurate answer

```

possible for each question, ensuring high-quality and useful responses for the user.  
.\n"

"Now, let's analyze the following text:\n"

f"Text: {chunk}\nQuery: Does this text mention any physical assets, locations or ownerships? Does the text mention what commodity the physical asset is being used for?\n"

"If yes, you must specify them in the following format:\n"

"physical assets: [ ]\nlocations: [ ]\nownerships: [ ]\ncommodities: [ ]\n"

"Additionally, identify the relationships between them, specifying the location of each physical asset, the ownership details, the commodity the physical asset is used for and the status of the physical asset."

"Format the relationships as follows:\nrelationships: [asset: '', location: '', ownership: '', commodity: '']."

)

### Role + instructional prompting

"You are a virtual assistant with advanced expertise in a broad spectrum of topics, equipped to utilize high-level critical thinking, cognitive skills, creativity, and innovation.\n"

"Your goal is to deliver the most straightforward and accurate answer possible for each question, ensuring high-quality and useful responses for the user.  
.\n"

"A physical asset is a tangible resource that a company owns and uses in the production of goods and services.

Examples of physical assets are facilities, equipment, infrastructure, etc. Ensure that a geographical location or region is never considered as an asset.  
.\n"

"A financial asset or other non-physical asset should never be included as a physical asset. Examples of financial assets include equity commitments, corporate facilities, accounts receivable, and short-term investments. Never include these in the list of physical assets.  
.\n"

"A commodity is what the physical asset is being used for. Examples include copper, gold, electricity, renewable energy, etc."

"Now, let's analyze the following text:\n"

f"Text: {chunk}\nQuery: Does this text mention any physical assets, locations or ownerships? Does the text mention what commodity the physical asset is being used for?\n"

"If yes, you must specify them in the following format:\n"

"physical assets: [ ]\nlocations: [ ]\nownerships: [ ]\ncommodities: [ ]\n"

"Additionally, identify the relationships between them, specifying the location of each physical asset, the ownership details, and the commodity the physical asset is used for."

"Format the relationships as follows:\nrelationships: [asset: '', location: '', ownership: '', commodity: '']. Do not output anything else."

)

## IRZ-CoT prompting

```
prompt_instruction = (  
    "You are a virtual  
    assistant with advanced  
    expertise in a broad  
    spectrum of topics,  
    equipped to utilize  
    high-level critical  
    thinking, cognitive  
    skills, creativity, and  
    innovation.\n"  
    "Your goal is to deliver  
    the most  
    straightforward and  
    accurate answer  
    possible for each  
    question, ensuring high  
    -quality and useful  
    responses for the user  
    .\n"  
    "A physical asset is a  
    tangible resource that  
    a company owns and uses  
    in the production of  
    goods and services.  
    Examples of physical  
    assets are facilities,  
    equipment,  
    infrastructure, etc.  
    Ensure that a  
    geographical location  
    or region is never  
    considered as an asset  
    .\n"  
    "A financial asset or  
    other non-physical  
    asset should never be  
    included as a physical  
    asset. Examples of  
    financial assets  
    include equity  
    commitments, corporate  
    facilities, accounts  
    receivable, and short-  
    term investments. Never  
    include these in the  
    list of physical assets  
    .\n"  
    "A commodity is what the  
    physical asset is being  
    used for. Examples  
    include copper, gold,  
    electricity, renewable  
    energy, etc."  
    "Now, let's analyze the  
    following text:\n"  
    f"Text: {chunk}\nQuery:  
    Let's think step-by-  
    step. Does this text  
    mention any physical  
    assets, locations or  
    ownerships? Does the  
    text mention what  
    commodity the physical  
    asset is being used for  
    ?\n"  
    "If yes, you must specify
```

```
    them in the following  
    format:\n"  
    "physical assets: [ ]\  
    nlocations: [ ]\  
    nownerships: [ ]\  
    ncommodities: [ ]\n"  
    "Additionally, identify  
    the relationships  
    between them,  
    specifying the location  
    of each physical asset  
    , the ownership details  
    , and the commodity the  
    physical asset is used  
    for."  
    "Format the relationships  
    as follows:\n    nrelationships: [asset:  
    ', location: ',  
    ownership: ',  
    commodity: '']. Do not  
    output anything else."  
    )
```

## Dynamic prompting

```
prompt_instruction = (  
    "You are a virtual  
    assistant with advanced  
    expertise in a broad  
    spectrum of topics,  
    equipped to utilize  
    high-level critical  
    thinking, cognitive  
    skills, creativity, and  
    innovation.\n"  
    "Your goal is to deliver  
    the most  
    straightforward and  
    accurate answer  
    possible for each  
    question, ensuring high  
    -quality and useful  
    responses for the user  
    .\n"  
    )  
  
    if contains_assets:  
        prompt_instruction += (  
            "A physical asset is a  
            tangible resource  
            that a company owns  
            and uses in the  
            production of goods  
            and services.  
            Examples of  
            physical assets are  
            facilities,  
            equipment,  
            infrastructure, etc  
            .\n"  
            "Ensure that a  
            geographical  
            location or region  
            is never considered
```



```

        as an asset.\n"
    "A financial asset or
    other non-physical
    asset should never
    be included as a
    physical asset.
    Examples of
    financial assets
    include equity
    commitments,
    corporate
    facilities,
    accounts receivable
    , and short-term
    investments. Never
    include these in
    the list of
    physical assets.\n"
)

if contains_commodities:
    prompt_instruction += (
        "A commodity is what
        the physical asset
        is being used for.
        Examples include
        copper, gold,
        electricity,
        renewable energy,
        etc.\n"
    )

if contains_locations:
    prompt_instruction += (
        "Always ensure that a
        geographical
        location or region
        is mentioned
        separately from the
        physical asset.\n"
    )

prompt_instruction += (
    f"Now, let's analyze the
    following text:\n"
    f"Text: {chunk}\nQuery:
    Let's think step-by-
    step. Does this text
    mention any physical
    assets, locations or
    ownerships? Does the
    text mention what
    commodity the physical
    asset is being used for
    ?\n"
    "If yes, you must specify
    them in the following
    format:\n"
    "physical assets: [ ]\
    nlocations: [ ]\
    nownerships: [ ]\
    ncommodities: []\n"
    "Additionally, identify
    the relationships
    between them,
    specifying the location
    of each physical asset
    , the ownership details

```

```

        , and the commodity the
        physical asset is used
        for."
    "Format the relationships
    as follows:\
    nrelationships: [asset:
    ', location: ',
    ownership: ',
    commodity: '']. Do not
    output anything else."
)

```

## Database Cleaning

You are an expert data cleaner.  
Your task is to clean and standardize the following text.  
You will be provided each cell value one by one with its respective column name. Apply the following cleaning steps:

- Standardize entries in the commodity column to have a consistent format. For example, "Silver, Gold, Lead , Zinc" should be the standard format for each commodity, separated by commas and no extra spaces.
- Ensure all entries in the status column are in a consistent format, removing redundant words or phrases.
- All entries should be in title case.
- Do not make any changes to the 'Countries' column.
- In the 'commodity' column, if chemical symbols are given, change these to the element name corresponding to the chemical symbol.
- In all the columns, ensure each entry is properly formatted without redundant commas and extra spaces. For example, "ExxonMobil" should not be separated by extra commas.
- Remove any leading or trailing spaces in all columns.
- All individual commodities should be separated by a comma.
- The 'location' column should only consist of geographical regions and locations.
- A physical asset is a tangible resource that a company owns and uses in the production of goods and services. Examples of physical assets are

- facilities, equipment, infrastructure, etc. If there are any entries in the physical asset column that do not fit the description of a physical asset, put N/A in the corresponding cell.
- A commodity is what the physical asset is being used for. If there are any entries in the commodity column that do not fit the description of a commodity, put N/A next to the word in brackets.
  - Ensure that there are no repetitions or redundant entries in any of the cells.
  - If any cell has 'not specified', it should be empty.
  - All cells should have standardized entries.

Process the following text according to these instructions . Return only the new cleaned cell value, nothing else.

### Country Extraction Prompt

You are an expert in geographical locations. Given the location information provided, identify the countries mentioned in the location . Return the list of countries separated by commas. If no country is mentioned, return "N/A".

### Modified prompt for improvement module

"You are a virtual assistant with advanced expertise in a broad spectrum of topics, equipped to utilize high-level critical thinking, cognitive skills, creativity, and innovation.\n"

"Your goal is to deliver the most straightforward and accurate answer possible for each question, ensuring high-quality and useful responses for the user.\n"

"A physical asset is a tangible resource that a company owns in a location and uses in the production of goods and services. Examples of physical assets are all examples of 'Plant' in the tables ( Wateree, Greenville and Colonial Trail West are all physical assets).\n"

"A financial asset or other non-physical asset should never be included as a physical asset. Examples of financial assets include equity commitments, corporate facilities, accounts receivable, and short-term investments. Never include these in the list of physical assets.\n"

"A commodity is what the physical asset is being used for. The status of a physical asset gives information on whether the asset is operational, under construction or in end-of-life."

"Now, let's analyze the following text:\n"

f"Text: {text}\nQuery: Let's think step-by-step. Does this text mention any physical assets, locations or ownerships ? Does the text mention what commodity the physical asset is being used for?\n"

"Does the text mention the status of the physical asset? Examples of status include whether the asset is operational, under construction or in end-of-life."

"If yes, you must specify them in the following format:\n"

"physical assets: [ ]\nlocations: [ ]\nownerships: [ ]\ncommodities: [ ]\nstatus: [ ]\n"

"Additionally, identify all the relationships between them, specifying the location of each physical asset"

```

, the ownership details
, the commodity the
physical asset is used
for and the status of
the physical asset. Do
not leave out any
relationships. "
"Format the relationships
as follows:\
nrelationships: [asset:
'', location: '',
ownership: '',
commodity: '', status:
'']. Do not output
anything else."

```

#### A.14 Standardised Characters in Foundational Data Cleaning & Standardisation

In this section, we detail the standardisation rules applied during the Foundational Data Cleaning & Standardisation phase in our database cleaning process. By enforcing these cleaning and standardisation rules, the data becomes more reliable and easier to analyse, reducing potential errors caused by inconsistent naming and formatting practices. This standardisation forms the foundation for subsequent cleaning steps within the study.

- **Characters Standardised:**

- \ (Backslashes): Removed from text fields.
- ' (Single quotes): Removed from text fields.
- " (Double quotes): Removed from text fields.
- **Extra spaces:**
  - \* Leading and trailing spaces: Trimmed from text entries.
  - \* Multiple consecutive spaces: Condensed to a single space.
- **Commas (,):** Ensured proper spacing after commas by replacing them with ' , ' (comma followed by a space).

- **Names Standardised:**

- **Company Names in ownership Field:**
  - \* All the following aliases and variations are standardized to **"Newmont Corporation"**:
    - "Company"
    - "company"
    - "The company"

- "the company"
- "The Company"
- "Company's"
- "the Company"
- "we"
- "NEWMONT CORPORATION"
- "Newmont's ownership or economic interest"
- "Company owns or controls land"
- "Newmont"
- "Newmont (majority)"
- "Newmont Corporation (formerly)"
- "100% owned by the Company"
- "Newmont Stockholders"
- "100% by Newmont"
- "Company owned"
- "Company's"

- **USA Variants in location Field:**

- \* All the following variants are standardised to **"USA"**:
  - "United States of America"
  - "United States"
  - "USA"
  - "US"
  - "USAA"
  - "USAA."
  - "U.S."
  - "U.S.A."

#### A.15 LSEG Database Numerical Results



In this section, we present the complete numerical results from stage 1 of the validation, showing six performance scores for each company.

Table 9: Complete Performance Scores for Companies across the following metrics: Partial Match Score, Jaccard Similarity, Cosine Similarity, Dice Similarity Coefficient, Normalised Levenshtein Distance, and Hits@5.


Sector	Company	Entity	Partial Match	Jaccard Similarity	Cosine Similarity	Dice Similarity	Norm. Levenshtein	Hits@5
Mining	AA	Physical Asset	0.95	0.7833	0.8591	0.8333	0.7667	0.7143
		Ownership	0.93	0.2139	0.3609	0.3389	0.3658	0.7143
		Commodity	0.52	0.0000	0.0000	0.0000	0.1839	0.7143
		Countries	0.100	0.8333	0.8963	0.8333	0.8639	0.7143
	SCCO	Physical Asset	0.81	0.6389	0.7416	0.7333	0.6349	0.5455
		Ownership	1.00	0.3000	0.5375	0.4127	0.5590	0.5455
		Commodity	0.68	0.0000	0.5244	0.0000	0.1966	0.5455
		Countries	0.82	0.6667	0.6667	0.6667	0.7222	0.5455
	FCX	Physical Asset	0.88	0.5667	0.7014	0.6672	0.5148	0.8462
		Ownership	0.93	0.1667	0.7063	0.2222	0.5020	0.8462
		Commodity	0.87	0.0000	0.3668	0.0000	0.2114	0.7692
		Countries	1.00	1.0000	1.0000	1.0000	1.0000	0.8462
	NEM	Physical Asset	0.83	0.8462	0.9083	0.8949	0.8531	0.6486
		Ownership	0.89	0.6520	0.7891	0.7090	0.7747	0.6216
		Commodity	0.83	0.1154	0.5488	0.1282	0.3698	0.5946
		Countries	0.79	0.7692	0.7692	0.7692	0.8225	0.6486
	HL	Physical Asset	0.82	0.7955	0.8608	0.8515	0.7738	0.6667
		Ownership	0.97	0.5000	0.6667	0.6667	0.7500	0.6667
		Commodity	0.88	0.0000	0.5776	0.0000	0.3281	0.6111
		Countries	0.62	0.1818	0.1818	0.1818	0.3566	0.6667
Oil & Gas	CVX	Physical Asset	0.72	0.3417	0.5084	0.5016	0.4292	0.6250
		Ownership	1.00	1.0000	1.0000	1.0000	1.0000	0.6250
		Commodity	0.36	0.0000	0.4472	0.0000	0.2857	0.1250
		Countries	0.53	0.0000	0.0000	0.0000	0.2000	0.5000
	XOM	Physical Asset	0.71	0.2418	0.4341	0.3840	0.3808	0.6364
		Ownership	0.94	0.4286	0.5469	0.5238	0.5276	0.6364
		Commodity	0.35	0.0000	0.0000	0.0000	0.1852	0.0000
		Countries	0.63	0.1429	0.1667	0.1429	0.3127	0.5455
	MPC	Physical Asset	0.74	0.3026	0.4935	0.4556	0.3552	0.7895
		Ownership	0.95	0.6083	0.7457	0.6925	0.6947	0.7368
		Commodity	0.49	0.1036	0.4149	0.1693	0.1907	0.5789
		Countries	0.43	0.4722	0.6855	0.6389	0.2493	0.7895
	COP	Physical Asset	0.76	0.2000	0.3536	0.3333	0.2500	1.0
		Ownership	1.00	0.3333	0.5774	0.5000	0.4118	1.0
		Commodity	1.00	0.3333	0.8944	0.5000	0.1667	1.0
		Countries	0.67	0.5000	0.7071	0.6667	0.2727	1.0
Utilities	D	Physical Asset	0.79	0.2532	0.6368	0.3961	0.3788	0.2778
		Ownership	0.85	0.2740	0.4884	0.4015	0.4966	0.3333
		Commodity	0.75	0.2292	0.3836	0.3125	0.2228	0.1667
		Countries	1.00	0.9375	0.9634	0.9583	0.8984	0.2778
	ED	Physical Asset	0.64	0.3500	0.6400	0.5000	0.4602	0.5000
		Ownership	1.00	0.5000	0.7071	0.6667	0.4637	0.5000
		Commodity	0.40	0.0000	0.0000	0.0000	0.1364	0.0000
		Countries	0.67	0.7500	0.8500	0.8333	0.6100	0.5000
	DUK	Physical Asset	0.74	0.3131	0.5573	0.4683	0.5001	0.6000
		Ownership	0.98	0.4667	0.7953	0.6286	0.5889	0.6000
		Commodity	0.41	0.0000	0.0000	0.0000	0.1051	0.0000
		Countries	1.00	0.5417	0.7155	0.6667	0.3789	0.6000
	EXC	Physical Asset	0.93	0.3778	0.7145	0.5367	0.5472	0.7333
		Ownership	0.71	0.2867	0.4354	0.4333	0.4136	0.6667
		Commodity	0.40	0.4000	0.6285	0.5333	0.3520	0.5333
		Countries	0.49	0.9000	1.0000	0.9000	1.0000	0.6667

## A.16 Dashboard User Interface (UI)


In this section, we present snippets of dashboard user interface (UI) for Alcoa Corporation (AA) as an example of the functionality and design of the visualisation tool developed for this project. The dashboard provides an intuitive and interactive platform to explore the relationships between key entities within the database.




Commodity vs. Ownership Heatmap




Commodities by Ownership



Commodities by Country



Ownership vs. Commodity vs. Countries Bubble Chart



Asset Distribution by Country and Ownership Treemap

