AI Agents for Document Automation and Negotiation Between Buyers and Sellers

Muhammed Abdurrahman Ismath*
SoukAI
abdurrahman@souk-ai.co.uk

Alaa Dania Adimi
SoukAI
dania@souk-ai.co.uk

February 28, 2025

Abstract

The oil and gas trading industry faces protracted deal cycles due to labor-intensive document handling and prolonged negotiations. This research proposal introduces an AI agent-based workflow to streamline document automation and broker negotiations between buyers and sellers. The core is an AI Broker Document Handling Pipeline that automates the extraction and structuring of trade documents and mediates communication between parties. By leveraging optical character recognition (OCR) and large language models (LLMs) for intelligent document processing and an AI broker agent to conduct negotiations with built-in guardrails, the system aims to dramatically compress deal timelines (from months to days) while maintaining trust and compliance. We outline the problem, objectives, methodological framework, and evaluation plan for this approach. The expected impact is a significant reduction in time-to-deal closure and improved efficiency in oil and gas transactions.

1 Introduction

1.1 Average Transaction Duration in Physical Oil Gas Trades

Traditional physical oil and gas transactions are notably slow and paperwork-intensive. Deals often rely on *manual documents and processes, which can stretch the timeline of a trade over multiple days or even weeks. For example, the negotiation and exchange of letters of credit and bills of lading – critical documents in a physical oil/gas trade – is typically a protracted affair. Industry case studies indicate that **finalizing trade documents and settlement can take “days or weeks” under manual systems [Roundtable: Technology and trade, the story so far — Global Trade Review (GTR)]. Even once a cargo is loaded, payment is not immediate; standard contract terms often give buyers around 30 days after loading to make payment, usually via bank-issued letters of credit Banks withdraw letters of credit for Russian oil, tighter sanctions feared. In practice, this means that a single oil shipment might only be paid and settled weeks after physical delivery. The reliance on couriers and email for document transfer contributes to delays – it is not uncommon that a vessel arrives at its destination before the paper documents do, underscoring the inefficiency. Indeed, a World Bank report on trade finance highlights that current settlement times “take days” rather than real-time. These delays tie up capital and expose parties to market and credit risk in the interim. By contrast, experiments with digital platforms have shown dramatic improvements: in one pilot, a process that normally took days or weeks was completed in a few hours using a blockchain-based system. This stark gap illustrates the average duration of traditional trades (on the order of several days) and the potential time savings if modernized.

1.2 Major Fraud Cases in Oil Gas Commodity Trading

The physical oil and gas trading sector has seen several high-profile fraud and market manipulation cases in recent years, exposing vulnerabilities in traditional trading practices. One of the most notorious was the collapse of Hin Leong Trading in 2020, a Singapore-based oil trader. Hin Leong’s founder admitted hiding $800 million in losses and was later accused of using forged documents to secure financing [Factbox: Commodities and energy trading firm scandals — Reuters]. The fallout led banks (including HSBC, ING, and others) into lengthy legal battles to recover some $3.5 billion in losses, with prosecutors alleging the company used “forged or fabricated documentation” to obtain credit [4 Billion Pieces of Paper Keep Global Trade Afloat but Fraud Fuels Push to Digitize]. Another case is ZenRock Commodities, also in Singapore (2020), where creditors like HSBC alleged “highly dishonest transactions” – the firm had over $600 million in outstanding debts and was found to have financed the same cargo multiple times, a fraud made possible by duplicate paper documents. Document forgery schemes have been a recurring theme: in one instance, employees of Coastal Oil Singapore were charged with creating fictitious sales contracts and invoices to defraud eight banks of over $340 million in commodity loans. Fraud isn’t limited to document forgery; it also includes unauthorized trading and market manipulation. For example, a trader at Mitsubishi Corp’s

Petro-Diamond unit in Singapore secretly built up huge derivative positions, resulting in a $320 million loss and the shutdown of the unit in 2019. Even major global firms have faced penalties: in 2024, Trafigura (one of the world's largest commodity traders) agreed to pay a $55 million fine to U.S. regulators for manipulating fuel oil benchmark prices using insider information. These cases underscore the scale of fraud and misconduct risks in the oil and gas trading sector – from forged shipping documents and financing scams to market price manipulation – often enabled by opaque, paper-based processes.

1.3 Global Inefficiencies: Delays, Losses due to Fraud, and Lack of Transparency

Physical oil and gas trading remains riddled with inefficiencies stemming from outdated practices. A striking indicator is the continued dominance of paper documentation in global trade. Paper documents still "rule" the $25 trillion global cargo trade, with an estimated 4 billion paper documents in circulation at any given time. These paper-based workflows are prone to errors, delays, and lack of visibility. Documents like bills of lading must be physically couriered between parties, which "frequently get lost" and can add huge delays to shipments and transactions. This slow, fragmented process not only slows down trade (increasing storage and financing costs due to delays) but also creates opportunities for fraud. According to the International Chamber of Commerce (ICC), at least 1% of all trade finance transactions (roughly $50 billion annually) are believed to be fraudulent. In the commodities sector alone, Bloomberg data shows that over the past decade, banks, traders and insurers have suffered at least $9 billion in losses due to falsified documents. Such frauds are enabled by the lack of transparency and verification in paper-based trading – criminals can present the same fake document (e.g. a bill of lading or warehouse receipt) to multiple banks to secure duplicate financing, because these institutions have no shared ledger to catch the duplication. As one anti-fraud expert noted, forging or duplicating paper trade documents is "the easiest type of fraud to commit" in this industry. Beyond fraud losses, inefficiencies also manifest as higher costs and missed opportunities. A McKinsey study cited by ICC found that moving to electronic bills of lading could unlock $40 billion in additional trade by reducing frictions, especially in emerging markets, and save global shipping lines around $6.5 billion per year in direct costs. Moreover, the lack of transparency in physical trading can lead to distrust and disputes. Opaque pricing and uneven information flow make it harder to reconcile transactions, often requiring lengthy manual reconciliations and audits. Recognizing these issues, major industry players and regulators have begun pushing for digitization and transparency. Yet, as of 2023, less than 2% of global trade transactions are conducted via digital systems – meaning the vast majority of trades still rely on century-old paper processes. This low adoption of digital tools highlights a significant global efficiency gap. In summary, the inefficiencies in physical oil and gas trading include protracted delays in transaction completion, substantial losses to fraud facilitated by document forgery, and a general lack of transparency that undermines trust and smooth operation. Industry reports from groups like the OECD and ICC, as well as investigations by Reuters and Bloomberg, all emphasize that modernizing these legacy processes (through electronic documents, data sharing, and better oversight) is crucial to reducing delays and fraud in the oil and gas commodity sector.

2 Methodology

Our proposed solution comprises a multi-agent system with two primary components: a **Document Handling Pipeline** and an **AI Broker agent**, supplemented by a matchmaking feature. The design focuses on end-to-end automation of the deal workflow, from initial document processing to final negotiation, with robust safety measures in place.

2.1 Document Handling Pipeline

Oil and gas deals involve numerous documents (contracts, letters of credit, technical specifications, etc.) that are often in PDF or scanned formats. The document-handling pipeline automates the ingestion and understanding of these materials as shown in Figure 1.

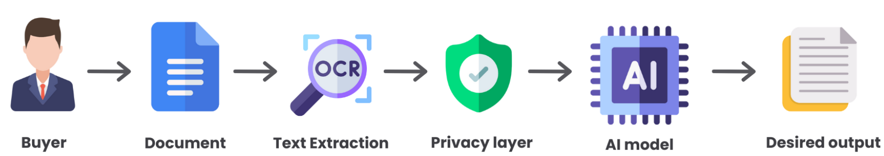

Figure 1: Documents handling pipeline

First, an OCR module extracts text from scanned documents or images with high accuracy. Next, a large language model is employed to interpret and structure the content. Modern LLMs are capable of understanding domain-specific text and extracting key information at an unparalleled level of comprehension [4]. The pipeline uses prompt-based LLM processing to identify critical fields (e.g., product details, quantities, pricing terms, delivery conditions) and convert unstructured text into a structured data format. It also generates concise summaries of lengthy documents, highlighting essential points and any anomalies. This LLM-driven approach to document analysis ensures that both parties have a clear, summarized understanding of all documents, drastically reducing the time needed for manual review. Importantly, a dedicated data privacy layer is integrated into the pipeline; this layer automatically redacts any Personally Identifiable Information (PII) before sending the document's content to the external language model, ensuring sensitive data remains protected and compliance with privacy regulations is maintained.

2.2 AI Broker for Negotiation

The AI Broker is a multi-agent system that mediates communication between the buyer and seller, effectively acting as an intermediary (see Figure 2). Once the documents are processed and both parties' positions (e.g. desired terms, constraints) are known, the AI Broker initiates a negotiation dialogue. It interfaces with each party (likely through a chat or messaging platform), presenting questions, offers, and counter-offers on behalf of the other party.

Crucially, the AI Broker ensures trust and transparency in several ways:

1. It maintains a separation layer so that buyers and sellers do not exchange identifying information or communicate outside the broker channel, preventing premature disclosure of sensitive details.

2. It keeps an auditable log of all interactions so that the negotiation process is transparent and traceable.

3. It adheres to predefined negotiation protocols and etiquette, ensuring that the dialogue remains professional and focused on the terms of the deal.

The AI Broker uses LLMs to understand and generate proposals in natural language. It is capable of reasoning about the trade-offs between different terms (e.g., price vs. delivery time) and can suggest fair compromises. If

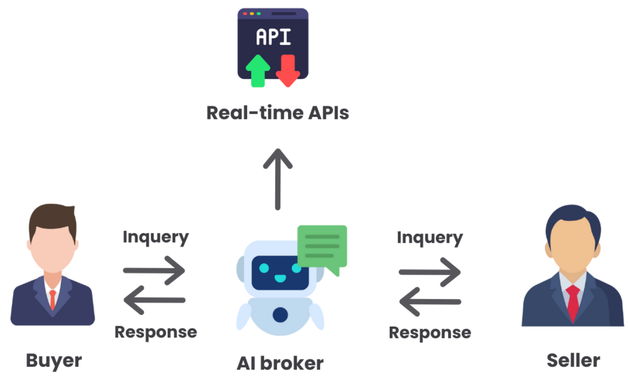

Figure 2: Documents handling pipeline

either party asks a question that can be answered from the documents (for example, "Does the quality certificate cover batch XYZ?"), the AI broker will consult the processed document data to provide an answer, behaving like an intelligent assistant with full context of the deal. Throughout the negotiation, the AI Broker leverages real-time pricing data via integrated APIs to ensure both parties have current market rates at their fingertips. It balances the objectives of both sides to move toward a mutually acceptable agreement. By handling the bulk of the communication—thus eliminating endless email threads and calls—the AI Broker frees human negotiators to step in for final approval or for decisions that fall outside its operational scope.

2.3 AI Safety Measures in the Broker

While the AI Broker brings efficiency, it is imperative to integrate strong AI safety measures to maintain reliability and confidentiality. LLM-based agents can sometimes produce inaccurate or inappropriate outputs if not properly controlled [1]. Our system, therefore, leverages guardrails [2]. Recent research emphasizes incorporating LLM guardrails to ensure responsible use and to prevent failure modes in deployed AI systems. In our broker, guardrails are put in place to check the AI's messages before they reach the user. These guardrails serve to mitigate risks such as accidental disclosure of private data, biased or toxic language, or legally non-compliant suggestions. In practice, the guardrail mechanism will intercept the AI's draft outputs and check for policy violations. For instance, if a seller's prompt inadvertently elicits buyer's confidential information, the guardrail will block or redact that content.

Privacy is a top priority. The AI Broker is designed with privacy layers to ensure that sensitive business information is protected. As mentioned above, the system employs privacy-preserving techniques so that the AI can learn negotiation strategies without directly exposing any private details. The AI broker will only share information that is approved and necessary for negotiation, and any personal identifiers or extraneous details are stripped out.

3 Evaluation

We will evaluate the proposed systems on both its technical performance and its business impact. Key metrics include:

1. **Document Extraction Accuracy:** Measuring how accurately the pipeline converts PDFs/scans into structured data. This will involve human and automatic evaluation using approaches such as LLM-as-judge [3].

2. **Negotiation Efficiency:** Quantifying the reduction in deal closure time. In pilot deployments, we will track the duration from initial contact to final agreement when using the AI systems compared to historical averages without AI. Our goal is to demonstrate a reduction from the scale of months to days in closing deals.

3. **Safety and Compliance:** Testing the AI Broker's guardrails by attempting various adversarial scenarios. For example, we will conduct prompt injection tests and ensure the broker consistently deflects them. Success is measured by zero occurrences of unauthorized information leakage or policy violations during simulations..

By measuring these aspects, we can iteratively improve the system. A successful evaluation would show that the AI agents handle documents as accurately as humans, negotiate deals significantly faster, and do so in a secure and trustworthy manner.

4 Contribution and Future Research

This research will contribute a novel framework for agentic automation in commodity trading. The AI Broker Document Handling Pipeline demonstrates how combining NLP capabilities with autonomous agent coordination can transform a traditionally slow, human-intensive domain. The immediate added value is practical: faster deal closures, lower overhead costs, and enhanced accuracy in compliance and documentation. For the oil and gas industry, this could mean more liquidity and agility in trading operations, as deals that once stalled due to paperwork can now progress quickly. Moreover, our approach can serve as a template for other industries involving heavy documentation and negotiation (such as finance, real estate, or insurance).

In terms of academic contributions, this project explores the intersection of multi-agent systems, natural language understanding, and human-agent interaction in a high-stakes, real-world scenario. We integrate state-of-the-art OCR and LLM techniques into an agent workflow and address challenges of trust and safety in automated negotiations. The findings from this work could inform future research in several areas. One avenue is improving the reasoning and strategy of the AI Broker: for example, incorporating reinforcement learning so the broker can learn optimal negotiation tactics over time (adapting strategies as seen in prior autonomous negotiating agents research). Another area is expanding the AI safety mechanisms – as AI agents take on more complex tasks, robust alignment with human values and legal constraints becomes even more crucial. Techniques like more advanced truthfulness checks, explainability of the AI's decisions, or cross-agent verification (where multiple AI agents cross-validate each other's outputs) could be investigated. Additionally, scaling the matchmaking and negotiation to support marketplaces with many participants and multi-deal optimization could be a future direction. Overall, the proposal not only aims to solve an immediate industry problem but also to pave the way for more intelligent and safe AI agent frameworks in business transactions.

5 Conclusion

In conclusion, the proposed AI agent framework stands to revolutionize the document-heavy and negotiation-intensive workflow of oil and gas trading. By automating document handling and introducing an AI broker to facilitate negotiations, we target a dramatic improvement in transaction speed and efficiency. The integration of safety guardrails and privacy features ensures that this increased autonomy does not come at the expense of security or trust. Early results will seek to validate that deals that once took months can indeed be concluded within days under this AI-mediated approach without sacrificing accuracy or fairness. This research thus promises to contribute both to industry practice and to the broader field of AI agents, demonstrating a tangible step towards smarter, safer autonomous negotiations.

References

[1] John Doe, Jane Smith, and Alex Lee. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *arXiv preprint arXiv:2305.11391*, 2023. 2023d.

[2] Yi Dong, Ronghui Mu, Gaojie Jin, Yi Qi, Jinwei Hu, Xingyu Zhao, Jie Meng, Wenjie Ruan, and Xiaowei Huang. Building guardrails for large language models. *arXiv preprint arXiv:2402.01822*, 2024.

[3] Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Yuanzhuo Wang, Wen Gao, Lionel Ni, and Jian Guo. A survey on llm-as-a-judge. *arXiv preprint arXiv:2305.11391*, 2023.

[4] Derong Xu, Wei Chen, Wenjun Peng, Chao Zhang, Tong Xu, Xiangyu Zhao, Xian Wu, Yefeng Zheng, and Enhong Chen. Large language models for generative information extraction: A survey. *arXiv preprint arXiv:2312.17617*, 2023.