

Harnessing the Power of AI for Commodities Markets

Energy, Agriculture, and Metals:
The AI Revolution Arrives for
Trading and the Back Office



Table of contents

01

Introduction

02

Defining key areas
of AI innovation

03

AI in practice
for commodities
markets with
ClearDox

04

Benefitting from
AI solutions by
ClearDox

05

Utilize AI in
commodities
markets now —
the right way





1

Introduction

When ChatGPT was launched at the end of November 2022, it seemingly began the “Age of AI” much as the launch of Sputnik in 1957 began the space race. And just like Sputnik, the resulting AI frenzy is loaded with misinformation, half-baked assessments, and incorrect conclusions that make it difficult to separate fact from fiction.

Take a step back from the headlines however, and a dual truth emerges: the roots of artificial intelligence go back for decades. And looking past the hype, there are a large number of significant developments that are already in use and fundamentally transforming business.

The challenge for the enterprise is to determine how to best take advantage of the advancements that are currently available rather than chase the latest headline. The objective is to match the technology to the task, not the other way around, and to do so requires both insight and patience.

With a little research it becomes clear that AI is already impacting business — and shows promise to deliver bigger results.

A case in point exists in the trading and the back office world of physical commodities. Advanced analytical technologies have been deployed to solve some of the most complex and intractable challenges. And new frontiers of AI tools are leading to dramatic increases in efficiency, risk management capabilities, and analytic insight.

In this white paper, we explore the universe of new AI tools and examine specific use cases from energy, agricultural, and metals markets that illustrate how these tools are making an impact both now and in the future. In addition to the challenges that must be addressed, we also describe the potential advancements that lie just beyond the horizon. Overall, we’ll learn that the effective utilization of AI is case dependent and that a combination of technical expertise and market experience is a requirement for optimal results.

2

Defining key areas of AI innovation

In any analysis, it's always important to begin from a point of common understanding by defining terms and setting the scope of conversation.

While consumer facing advancements like text generation with ChatGPT or PaLM, unique artwork through DALL-E or DeepArt, or original music composition from Aiva or Magenta Studio are powerful, potentially useful and intellectually intriguing, without a clear purpose and objective they can be more like toys than tools. Therefore, it helps to know the capabilities, strengths and weaknesses of each category of AI so any technology employed is the optimum choice for the task at hand.

Here are some of the key components ClearDox is utilizing for commodities trading and back office functions in energy, agriculture, and metals markets:



1. Optical character recognition (OCR)

OCR is a technology that converts images of text into a machine-readable format. It has been used for many years to digitize paper and electronic documents, making them easier to store, search, and edit. OCR works by scanning the image of the document and then using pattern recognition to identify the characters.

OCR by itself is only marginally useful, as the raw output is just a series of characters. But when combined with the position of the text on the document, it can produce structured output. For example, an invoice typically contains an invoice number, invoice due date and invoice amount due. By applying OCR and position, the information can be retrieved off the image and put into a structured format.

Document Viewer

Pages | Converted Document

Invoice

AmSpec LLC
1249 South River Rd.
Suite 204
Cranbury, NJ 08512
United States

Invoice Number: 302-1854554
Invoice Date: 17-Jan-2022
Payment Terms: As per contractual agreement
Supporting Documentation: [Click here to download](#)

Bill To: International Petroleum Corporation
1001 28th Street
Ocala, Virginia
33129
United States

Contact Name: Kipdon Colton

Job Order | Job Type | Job Location | Product Grade (volume) | Inspection Object | Job Finish Date | Scheduler | Office/Contact Details

Description	Qty	Units	Unit Price	%	%	Total
1	1	Sam	1,350.0000	100.00	75.00	1,012.50
2	1	Sam	204.0000	100.00	75.00	153.00
3	1	ea	302.0000	100.00	75.00	226.50
4	1	ea	0.0000	100.00	75.00	0.00
5	1	ea	243.0000	100.00	75.00	182.25
6	1	ea	243.0000	100.00	75.00	182.25
7	1	ea	243.0000	100.00	75.00	182.25
8	1	ea	0.0000	100.00	75.00	0.00
9	1	ea	195.0000	100.00	75.00	146.25
10	1	ea	243.0000	100.00	75.00	182.25
11	1	ea	475.0000	100.00	75.00	356.25
12	1	ea	396.0000	100.00	75.00	297.00
13	1	ea	599.0000	100.00	75.00	449.25
14	1	ea	1.51	100.00	75.00	1.13
15	1	ea	443.0000	100.00	75.00	332.25
16	1	ea	428.0000	100.00	75.00	321.00
17	1	ea	428.0000	100.00	75.00	321.00
18	1	ea	428.0000	100.00	75.00	321.00
19	1	ea	1,286.0000	100.00	75.00	964.50
20	1	ea	794.0000	100.00	75.00	595.50
21	1	ea	145.0000	100.00	75.00	108.75
22	1	ea	195.0000	100.00	75.00	146.25
23	1	ea	894.0000	100.00	75.00	670.50
24	1	ea	877.0000	100.00	75.00	657.75
25	1	ea	396.0000	100.00	75.00	297.00
26	1	ea	596.0000	100.00	75.00	447.00
27	1	ea	399.0000	100.00	75.00	299.25
28	1	ea	888.0000	100.00	75.00	666.00
29	1	ea	360.0000	100.00	75.00	270.00

Document Parser

Lines | Fields | Tables | Custom Script | Settings

Page-1 | Type text to filter...

Invoice Number: 302-1854554 | **Invoice Date:** 17-Jan-2022

Payment Terms: As per contractual agreement | **Reference Number:** 3742876-64

Supporting Documentation: Click here to download

Client Code: 458456 | **PO Number:**

OCR captures data from specific locations on forms and populates a table.

OCR was first introduced nearly 50 years ago, so it is well developed. It has excellent capabilities for handling image-based documents while being compatible with many languages.

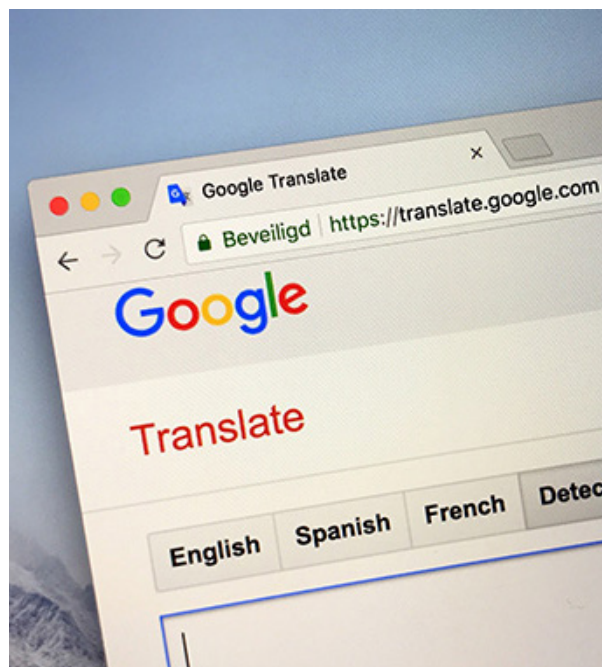
However, OCR's usefulness is limited by the fact that it relies on the position of text within a document: templates need to be built for each type of document. This can be onerous once more unique documents are encountered because each table requires custom scripting. These tables can be brittle as well: A change in format for an input renders an existing table useless.

All in all, OCR is a foundational technology that is the best choice when the information being gathered for analysis is highly structured, consistent, and not subject to changes in format.

2. Natural Language Processing (NLP)

NLP focuses on the interaction between computers and humans through natural language. It enables machines to understand, interpret, and generate human language, facilitating tasks like translation, sentiment analysis, speech recognition, and information extraction. It's a multidisciplinary field, drawing from computer science, artificial intelligence, and linguistics.

Some well-known types of NLP (and their classification) include Google Translate (machine translation), Siri and Alexa (speech recognition), and chatbots or virtual assistants (a combination of both).



“In my opinion, one of the hardest problems in machine learning is figuring out what problems I can solve with machine learning.

There’s still a gap in the understanding.”

– Mikey Shulman, lecturer at MIT Sloan and Co-founder, Suno AI¹

At ClearDox, recent solutions have utilized Named Entity Recognition (NER), which is a type of information extraction that identifies predefined categories of “named entities” in a body of text. For example, entity labels can be defined to extract the invoice number, invoice due date and invoice amount due from an invoice document. This requires a set of documents to be trained through an annotation process that provides the input into the NLP/NER model.

NER is extremely useful because it is readily capable of handling documents that have a great deal of unnecessary or redundant text and, unlike OCR, it scales very well to handle the addition of new documents for input. In all, it is a targeted solution that solves a specific language problem very well.

On the other hand, NER does not easily handle tabular data and it can also be more expensive to build and maintain, both because it requires up front training and because model monitoring and uptraining are needed to ensure smooth operation and output.

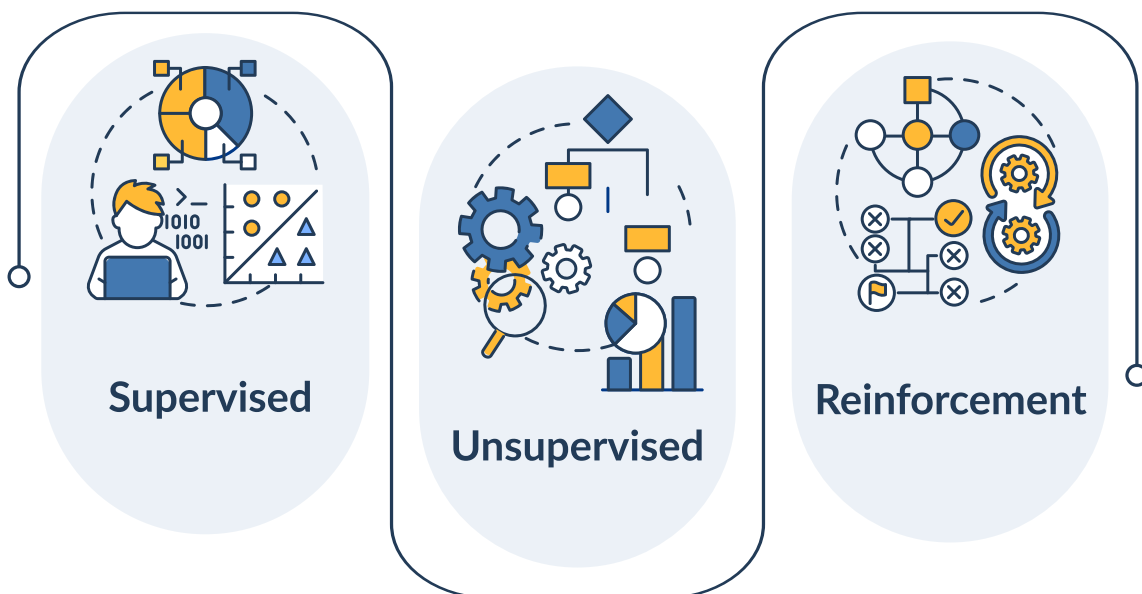
1. <https://mitsloan.mit.edu/ideas-made-to-matter/machine-learning-explained>

3. Machine Learning (ML)

ML is a subfield of AI, although the terms are sometimes conflated. ML is broadly defined as the capability of a machine to imitate intelligent human behavior. ML systems are used to perform complex tasks in a way that is similar to how humans solve problems. The goal is to give computers the ability to “learn” without explicitly being programmed.

There are three subcategories of ML:

- **Supervised ML** models are trained with labeled data sets, which allow the models to learn and grow more accurate over time. For example, an algorithm would be trained with examples of invoices or bills of lading, all labeled by humans. The machine would use those examples to learn how to identify the same information in other documents all on its own. Supervised machine learning is the most common type used today.
- In **unsupervised ML** a program looks for patterns in unlabeled data. Unsupervised machine learning can find patterns or trends that people aren’t explicitly looking for. For example, an unsupervised machine learning program could look through commodities transaction data and identify changes in consumption or sales patterns over time.
- **Reinforcement ML** trains machines through trial and error to take the best action by establishing a reward system. Reinforcement learning can train models to play games or train autonomous vehicles to drive by telling the machine when it made the right decisions, which helps it learn over time what actions it should take.



In one instance, the team at ClearDox has utilized a form of supervised ML called Form Parser. It uses ML to identify form fields and tables on a document.

Because forms contain so much valuable information that is often hidden or obscured, form extraction is an alluring field. Form Parser is a general-purpose form extraction tool that attempts to find relevant data in any type of input document, not just common items like receipts or invoices, giving it greater potential utility compared to OCR or NLP. It is also more powerful because fields and tables are automatically found, and are not dependent on text position within a source document.

Weaknesses include a requirement for substantial post processing to map fields and tables into a structured schema and the fact that, as the name implies, Form Parser works best on form-based layouts.

4. Generative AI

In a word, Generative AI is what all the AI noise is about these days. The category includes ChatGPT, the consumer-facing AI product that was introduced by OpenAI in November 2022 and started the “Sputnik” moment for AI. In addition to the obvious benefit of putting an intuitive and easy to use AI tool in more hands, the technological approaches are a further development of more powerful and, therefore, useful AI tools.

The technological underpinning of Generative AI is a large language model (LLM). An LLM is a type of machine learning model that can perform a



variety of natural language processing tasks such as generating and classifying text, answering questions in a conversational manner, and translating text from one language to another.

The label “large” refers to the number of values (parameters) the model can autonomously change as it operates and “learns”. It is possible for LLMs to have hundreds of billions of parameters that are trained over an extremely diverse set of content.

LLMs are an advanced manifestation of AI and have a number of benefits. They are usually trained on text artifacts that number in the billions, so they are very broad and deep, making them well suited to performing complex document analysis. They also require little or no upfront training so they can often be quickly deployed and utilized.

There are, however, a number of drawbacks and weaknesses as well. LLM output is unstructured, putting a premium on post-generation analysis, and Generative AI can require complex prompting to generate consistently meaningful results. Also, security issues are of special concern when sensitive, internal data is mingled with public data in an open source or other commercial model.

5. Graph Database

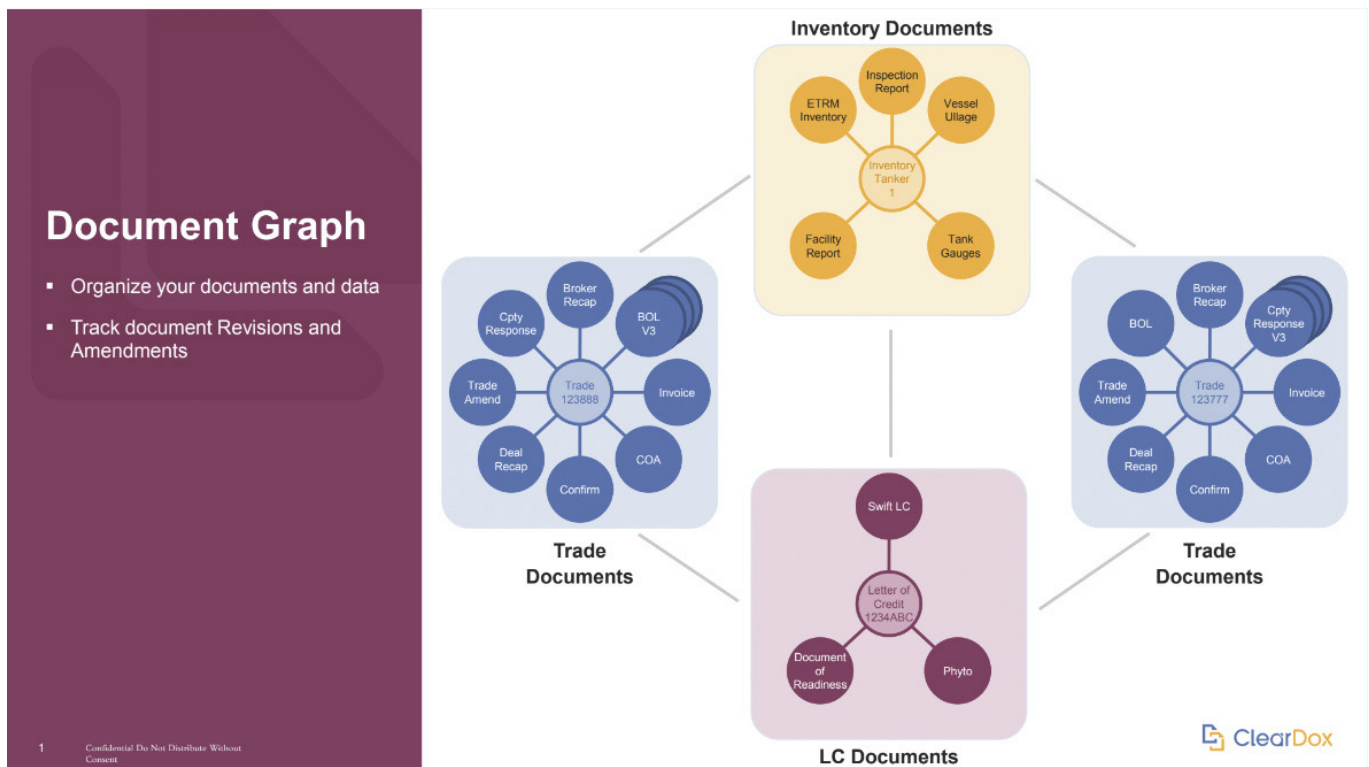
A graph database is not AI per se, but it is a tool that is instrumental in supporting AI efforts and making them more effective. A graph database excels at showing how data is interlinked, making it possible to recognize relationships within the data that cannot be communicated using a traditional, tabular SQL approach.

This is the primary technology behind FaceBook and LinkedIn and allows for content to be associated with other content through named associations or "edges". For example, User A (node) likes (edge) User B (node).

When this is extrapolated to a very large data set, the number of interconnections — and the insights they reveal — can expand exponentially.

Graph databases are especially effective in instances when it is important to recognize and define complex patterns quickly. In this way, a graph database is an excellent tool for enhancing artificial intelligence by revealing and describing the context of relationships in data.

As we grow our products and modules, ClearDox is targeting graph database technology as a means of building a Document graph that stores relationships across each document and the data within. With this technology in place, the depth of insight increases, as associations are more quickly and completely discovered and described.





3

AI in practice for commodities markets with ClearDox

When it comes to any technology implementation, the devil is in the details. Solutions that look good on paper may not be the best fit for a challenge. And even when they are, there is still a good deal of work to be done to adapt and plan so optimum results are obtained.

The mission at ClearDox is to digitize data from contracts, invoices, trading confirmations, bills of lading, freight receipts and more, so all relevant information is extracted and available for use for risk analysis, trading insights, regulatory reporting, fraud detection ... the list is virtually endless.

The sources of data vary widely, from tabular invoices to text-based contracts, requiring a variety of data extraction tools.

Here's how ClearDox utilizes these technologies:

OCR

OCR is often a foundation for any image-to-text digitization. ClearDox utilizes OCR but does not rely on text coordinates.

Instead, ClearDox defines a proprietary parsing language that uses anchor labels and relative positions from the anchor label. For example, invoice number is captured to the right of the text "Invoice No." This approach allows documents to be processed very quickly.

But, the solution does not scale well, as it still requires a template per document. OCR is best suited for uses where the inputs are highly structured and the documents are homogenous, with little expected format change or text complexity. OCR is rarely used by itself. Instead, it is a preprocessing step for feeding text from an image into a ML Form Parser, NLP model or LLM.



ClearDox uses three different OCR engines depending on the document to be processed. Each has their own advantages: It's a matter of selecting the most effective and efficient model for the task at hand:

- The basic OCR engine digitizes Latin languages (English, Spanish, French) swiftly in a cost-effective way and provides basic positional information.
- The more advanced OCR engine handles seven languages and provides deep positional information, allowing for rich post processing.
- The most advanced OCR engine handles 42 different languages, including Chinese and Arabic, and also supplies deep positional information for greater insight.

NLP

NLP is best suited to documents that are primarily text and is particularly strong in handling source material containing small nuggets of relevant information within a sea of definitions, clauses,

and other terms. Relevant data points are labeled within the text and trained to produce an NLP model. Hundreds of variations of the text are required to produce a robust model that can handle real-world language.

This is particularly helpful for oil contracts that tend to be very long but contain a lot of relevant information that has a great deal of specificity and variety and appears in many, sometimes multiple, locations in the documents. NLP is adept at capturing all of this information as a prelude to corroborating thoroughness and consistency in terms.

ML

The ML Form Parser is a generic model that recognizes form fields and tables on a document. This is ideal for any form-based document such as invoices, tax forms, bill of lading and inspection reports. Since the label names and table headers names are not known, significant post processing logic must be incorporated to map the data into a structured schema.



Generative AI

Generative AI using LLMs provides a significant advancement in document comprehension. This is ideal for comparing documents, searching documents, or interpreting the language within documents.

For example, when preparing a Letter of Credit presentation to a bank for financing, each document within the presentation package must be consistent and conditions met. LLMs can be used to compare the information in the Inspection Certificate, Bill of Lading and other supporting documents against the Letter of Credit to confirm consistency. The Additional Conditions section of the letter of credit can also be interpreted to ensure all documents are present and each condition is met, e.g., product is within 10 % of expected delivery volume, delivery is not from a restricted port, etc. In short, an LLM can open vast new horizons of insight.

Graph Database

The Document Graph is complementary to the AI technologies and is required for a company to contextualize data. For example, digitizing ten invoices or ten Bills of Lading documents is only marginally helpful to process the documents. However, knowing which contract, trade or letter of credit each document is associated with allows for straight through processing and auditability. The graph is used to store these relationships, building the complete data picture around a company's process. ClearDox is looking to leverage these graph technologies to track documents through the entire trade life cycle.

4

Benefitting from AI solutions by ClearDox

As commodity professionals with deep technology expertise, ClearDox provides data and analytics solutions that unlock information that is hidden within an organization on the way to delivering superior results for the enterprise. Some of the key tenets of the ClearDox approach are:



➤ Focus on data

It is axiomatic — but often overlooked — that you can't have good models if you don't have good data. At ClearDox, we recognize that sourcing, normalizing, and maintaining data is absolutely essential and we communicate and collaborate with our partners to make sure that the data landscape is always optimized.

➤ Results oriented approach

In any engagement, gaps will become apparent as work commences. These challenges will be solved in a step-by-step process every step along the way as functionality is added and expanded. Business objectives are always top of mind with ClearDox.

➤ Partner to manage complexity

As the world has become more and more complex it has become increasingly obvious that the old "build versus buy" question has largely been settled: Cooperation and partnership is essential in nearly all areas of an enterprise. In addition to the deep expertise that ClearDox brings to the table, our technical expertise is something you won't need to build or maintain on your own.

➤ The endpoint benefit is resiliency

In all ways, a partnership with ClearDox will increase your operational resilience. For example, as a new model learns, it will help you to quickly respond if there is an issue with a shipper, supplier, or other partner. Enhanced capabilities also help to increase cadence and capacity, accelerating how quickly you bring in new suppliers, vendors, or other parties because your system is more responsive and resilient.



5

Utilize AI in commodities markets now — the right way

Far too often, the emergence of a new technology leads to a “hammer looking for a nail” situation where business and technology leaders take a backwards approach, starting with a technology and then looking for a problem or challenge to apply it to.

That is certainly a danger right now with AI where a fear of missing out (FOMO) is likely to lead to suboptimal outcomes because the technology comes first and the issue at hand is a secondary consideration.

Fortunately, there’s a lot of light hiding behind all this noise..

Hype-cycles don’t last forever but they often lead to remarkable breakthroughs. Just as the launch of Sputnik in 1957 is directly related to the U.S. landing a man on the moon in 1969, the current frenzy and interest in the likes of ChatGPT signals the arrival of a new age that will solve many seemingly intractable problems and then go further, creating new opportunities and answering questions that have yet to be posed.

The realm of AI has deep roots and, at the same time, has the potential to grow in terms of both scope and scale as new developments lead to ever greater insights and the innovation cycle continues to feed on itself.

To avoid falling behind, the best time for an enterprise to get started is now.

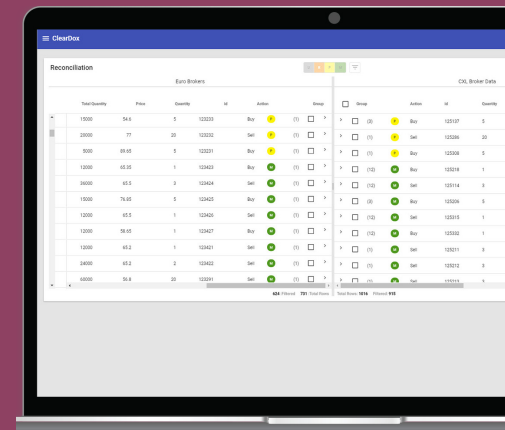


However, as discussed above, the best approach is to focus on the problems first and then look to find the best AI technology or technologies to apply to meet the challenges. This is the approach we have taken at ClearDox, applying our deep commodities market experience to analyzing challenges, and then building and deploying the most effective solutions to meet those challenges.

Combining experience with technology leads to the best possible outcomes. And when it comes to commodities-focused AI, we have the best of both.

To learn more about how ClearDox can revolutionize your commodity trading operations, book a 15-minute discovery meeting with our team. Take the first step towards enhancing efficiency and staying ahead in the fast-paced commodity trading landscape by leveraging ClearDox' deep experience and cutting-edge AI technologies.

BOOK A 15-MINUTE DISCOVERY CALL



About ClearDox LLC

ClearDox helps commodities companies turn manual processes into a competitive advantage. The ClearDox solution reduces costs, mitigates risk, and improves productivity by automating processes related to invoice processing, trade confirmations, inventory management, movement actualization and more. ClearDox, originally developed to solve data entry and reconciliation challenges for a global commodities merchant, classifies, extracts, and reconciles data using AI, ML and NLP before integrating it into downstream systems. An out-of-the-box solution that's easy to get up and running — and even easier to use — ClearDox has been trusted by companies including Gulf Oil, PBF Energy and Freepoint Commodities to process hundreds of thousands of documents since 2018. For more information visit cleardox.com.