




AI Opportunities Unlocked

(Generative) AI and its Application in the Oil and Gas Sector

Executive Summary

*This report explores the potential of **Generative AI within the oil and gas industry**. While offering significant opportunities for enhancing efficiency and sustainability, the successful implementation of Generative AI requires a nuanced understanding of its **capabilities and limitations**.*

Key findings include the ability of Generative AI to unlock valuable insights from unstructured data, optimize operations through predictive maintenance and fault diagnosis, and enhance safety protocols through real-time anomaly detection. The emergence of AI agents, capable of autonomous action and continuous learning, holds significant potential for optimizing operations and improving decision-making.

However, challenges include the potential for hallucinations and inaccuracies, data privacy and security concerns, and the risk of bias.

Successful implementation requires a phased approach, prioritizing data quality, robust security measures, and a thorough understanding of the technology's limitations. By carefully considering these factors, the oil and gas industry can leverage Generative AI to drive innovation, enhance efficiency, and improve overall performance.

Introduction

Author and futurist Stan Davis said: “when the infrastructure shifts, everything rumbles”. It’s a quote pertinent to the advances in Generative AI and its impact on almost every industry.

The Oil and Gas industry is no exception. This report is a guide to the art of the possible now with a view to the future to understand where the technology may lead. In a landscape that is continually evolving, even people working directly with AI find it challenging to plot the trajectory of the technology. This report isn’t a crystal ball, but I’ve attempted to provide a technologist’s view, a view grounded in reality, rather than a marketers view grounded in hype.

There is a growing consensus that we have hit “peak data” and the dramatic shifts in understanding are slowing. The efficiency of the foundational models however is increasing resulting in performance and reduced running costs. This has unlocked the application of the technology in more scenarios. For example, a reduction in latency enables more realistic real-time voice conversations allowing it to be used to handle voice calls like never before. The reduction in cost per token has made it commercially viable to process large volumes making tasks such as processing public consultations viable over the manual alternative.

Many Generative AI initiatives are incremental, simply enhancing current processes with AI. This mirrors the adoption of other technological leaps. When the electric motor was adopted in industry, it was simply used to replace existing drive shafts before the potential to revolutionise factory layouts was realised. What started as small improvements became a revolution. Everything rumbled.



The Evolving Landscape of the Oil and Gas Industry

The oil and gas industry is undergoing a profound transformation as it faces mounting pressures from various fronts. Operational efficiency is no longer just a competitive advantage; it has become a necessity to maintain profitability in an increasingly volatile market. Simultaneously, the industry is grappling with growing societal and regulatory demands for environmental sustainability. These forces are reshaping traditional approaches to exploration, production, and overall energy management.

Generative AI is emerging as a key enabler for innovation and progress. By leveraging advanced machine learning models, this technology offers the ability to unlock hidden insights from vast and complex data sets. For an industry that thrives on data—whether it be seismic surveys, reservoir simulations, or real-time sensor readings—this capability has the potential to enhance decision-making and uncover opportunities that were previously invisible.

Beyond data-driven insights, Generative AI also addresses inefficiencies in routine and repetitive tasks. By automating less valuable or time-intensive activities, such as document generation, report compilation, and even some aspects of compliance monitoring, organizations can redirect human capital toward higher-value strategic initiatives. The convergence of AI with traditional operations signals not just incremental progress but the possibility of transformative breakthroughs, paving the way for a more sustainable and efficient future for the industry.

It is crucial to acknowledge that not all AI solutions deliver a net benefit. Some implementations may incur higher costs than maintaining manual alternatives, underscoring the importance of carefully evaluating the cost-benefit ratio of any proposed solution. Security and data leakage remain paramount concerns, particularly within the oil and gas industry, where proprietary information and sensitive data are critical assets. Ensuring that data integrity and security protocols are robust is essential for mitigating risks associated with AI adoption.

Each oil and gas sector organisation is unique, with distinct operational challenges and objectives. As such, the success of AI applications depends on a tailored approach, grounded in a thorough understanding of the organization's specific context. While

generalized recommendations may inspire ideas, practical implementation requires a nuanced, case-by-case evaluation.

The following sections will explore areas where AI, particularly Generative AI, holds promise for the oil and gas industry. These insights aim to identify opportunities to enhance efficiency, innovation, and sustainability. Additionally, other areas of AI will be examined to provide a broader perspective on how this technology can address industry challenges.

Sections

- > Extracting Value from Unstructured Data
- > Operational
- > Safety in Operations
- > Rig and Machinery Monitoring
- > Security and Privacy
- > Data Privacy
- > IP Risks
- > Model Biases
- > Misinformation and Deepfakes
- > Emerging Threats
- > Mitigation



The Value of AI

Extracting Value from Unstructured Data

The oil and gas industry grapples with a wealth of unstructured data scattered across historical summaries, reports, and decision documents. This data, accumulated over time and encompassing complex subject matter, presents a significant challenge. Extracting valuable insights from this unstructured data often demands a substantial investment of effort, potentially exceeding the perceived value of the discoveries.

Traditionally, extracting structured insights from unstructured data has relied on heuristics and identifying patterns within language or signatures to discern trends. This broad-brush approach, however, often falls short. To truly unlock the value of this data, specialised translators are necessary. These tools would effectively "translate" tranches of data into a consistently readable format, enabling seamless processing and utilisation.

Generative AI enables us to take vast unstructured data and index it in such a way that it can be used against a foundational large language model (LLM) to provide answers and summarise.

Generative AI platforms are able to cite references to help researchers validate findings and delve further into the data found.

The unstructured data doesn't have to be processed or stored in a traditional data lake for Generative AI to use, although it simplifies the process. It is possible, akin to search indexing, to pull data from their existing stores. This might be from databases, file storage, documents, or any other digital form. Generative AI has the capability to read digital documents and spreadsheets and can make sense of images.

Operational

Operations stretch from office administration, plant processes, fault diagnosis and maintenance.

There are obvious applications of Generative AI in standard office administration contexts. Answering emails, creating schedules, generating reports and documents and day-to-day tasks.

Given the right prompting and templating, it is possible for Generative AI to generate the bulk of repetitive document creation. There are tools on the market that advertise that they can orchestrate the majority of tasks in this context. They can assist in writing emails, orchestrate arrangements and planning across applications and create templates and documents.

Examples of this are Microsoft's Co-Pilot and Google Gemini in Google Workspace. The reality, however, doesn't yet live up to the marketing. In addition to the mismatch in expectations, there are also security and privacy issues associated with the data access that these tools have. It is possible to mitigate these to some extent with guard rails and configuration, however, the limited data access then hinders its capability.

In the oil and gas sector, Generative AI also offers significant potential for optimising on-site operations. Downtime in oil and gas production is extremely costly. Generative AI can contribute to minimising downtime by assisting with equipment failure prediction, maintenance scheduling, and fault diagnosis. While not a standalone solution for maintenance scheduling, it can effectively support engineers by generating efficient plans and suggesting alternative approaches.

Predictive maintenance traditionally relies on machine learning models trained on substantial datasets. Generative AI can accelerate this process by augmenting existing data, particularly when historical data is sparse. Techniques like few-shot classification benefit from the synthetic data generation capabilities of Generative AI, enabling faster model training and more accurate predictions even with limited real-world examples.

A low effort with potentially large return is to use Generative AI to help field engineers with fault diagnosis and resolution. With a specifically trained and prompted Generative AI agent, field engineers could have an assistant on hand.

The AI could be given a photo and a short prompt, and a comprehensive explanation and support could be returned. Depending on the situation, it could delve into historical information, including past issues, maintenance, manuals and photos along with historical resolutions and provide the engineer with ideas and concepts to help them resolve the problem quickly. During the resolution, it could be asked additional questions such as how to diss-assemble a section or even more specifically, what torque settings to use when re-assembling.

In more advanced scenarios, the AI agent can even assist with logistics, such as identifying necessary replacement parts and initiating the ordering process through automated messaging.

The effectiveness of this AI agent will heavily depend on its training and the specific prompts provided. A generic AI model will likely have limited utility compared to one tailored to the unique needs and challenges of the oil and gas industry.

Safety in operations

Safety is paramount in industrial operations. Generative AI can significantly enhance safety protocols by generating comprehensive safety assessments based on specific work contexts. These assessments can include lists of required safety equipment and precautionary measures. Furthermore, AI-powered systems can proactively remind personnel to adhere to safety procedures before, during, and after each task.

Extending beyond textual analysis, Generative AI can leverage visual data to enhance safety. By analyzing existing CCTV footage, AI algorithms can identify anomalies and unusual activities within the work environment, triggering real-time alerts and notifications.

Computer vision technology offers even more advanced safety capabilities. Beyond human visual perception, computer vision systems can utilize alternative data sources, such as lidar, and advanced image processing techniques like spectral analysis and Eulerian amplification. These technologies enable the detection of subtle changes, such as vibrations or minute colour variations, that may indicate potential hazards invisible to the human eye.

It is crucial to carefully assess the specific safety challenges and opportunities within each industrial operation. Implementing AI-powered safety solutions should focus on practical applications that deliver tangible benefits, rather than merely adopting AI for novelty's sake.

Rig and Machinery Monitoring

Data can be collected either in batches or streamed in close to real-time to monitor rigs and plant machinery. Environmental monitoring can also be interleaved with this data to produce condition monitoring that goes beyond the direct data from couplings, pumps, and other rig machinery.

In cases such as mechanical seals, these function constantly but are difficult to monitor due to their inherent nature, especially internal variations. There are new technologies, specifically sensors, in the ultrasonics space that require no wiring or direct power and can accurately measure changes in internal structures.

We know of these sensors being used in the nuclear industry for internal monitoring, as well as within the "Christmas trees" in oil rigs to monitor the density of the material inside the valves. They are also used in mechanical seals.

It is possible to use this data for predictive maintenance, condition monitoring and throughput.

Predictive maintenance is the holy grail of systems like this, however, they require a lot of data and training to gain real value. Data on what not good looks like is important and has historically taken time to collect. To get value early, anomaly detection and other foundational models can be used to detect adverse trajectories of changes.

With rig and machinery monitoring in place, it is a perfect opportunity to apply Generative AI agents, an emerging combination of technologies that automate many tasks. This technology will be covered in a later section. Without data acquisition and connectivity, however, agent capability would either not be possible or be severely limited. Many of the outcomes build on subsequent technologies, highlighting the importance of having a strategy and plan for rolling out each layer so that value is added at each milestone and that the next milestone builds on the last.

Security & Privacy

Several security considerations arise when utilizing Generative AI, including data privacy, intellectual property (IP) risks, model biases, the spread of misinformation, deepfakes, and emerging threats such as prompt injection. While these threats are not entirely novel within the broader context of data and AI, the advent of readily accessible Generative AI technology facilitates the creation of highly convincing and manipulative communications with unprecedented speed and ease.

Data Privacy

Every piece of data shared with the providers of foundational large language models can and will be used to train subsequent models. This implies that if private information or Intellectual Property (IP) is used with these models, the data is effectively leaked. While OpenAI, through their APIs, claim to have mechanisms in place to prevent data leakage, this reliance on trust may be unacceptable for many companies.

When using tools like ChatGPT, all input data is ingested by the model. Therefore, comprehensive cybersecurity measures and practices are crucial to prevent the exploitation of AI models or the manipulation of their underlying data, which could lead to disruptions.

Integrated AI models, such as agents, present an increased risk due to the integration of inputs and outputs. These agents can take actions based on received inputs.

Since agents are exposed to external inputs, they become vulnerable to attacks. Furthermore, their ability to reflect changes in their environment can be exploited to gain control. Implementing robust security measures is paramount for all agent-style AI tools.

Data compliance is crucial, and all data usage must conform to regulatory standards such as GDPR and CCPA. Appropriate data anonymization and encryption are essential both in transit and at rest. Additionally, robust access controls must be implemented, aligned with the sensitivity of the data and the potential risks associated with its manipulation.

IP Risks

In some situations, especially when utilising foundational large language models, there is a risk that generated content may inadvertently incorporate copyrighted material. This may not always be readily apparent, potentially raising significant concerns.

Similarly, generated content may inadvertently incorporate or overlap with copyrighted or patented material, constituting intellectual property infringement. It is crucial to validate any content of intellectual value, just as with any internally generated ideas.

Model Biases

A common challenge with AI is that it can amplify biases in the data. This can be caused by the data having inherent biases and also the training and validation data sets being skewed. It is important to understand biases in data sets used for training and how the data is partitioned for training and validation to avoid excessive bias and overfitting.

Misinformation and Deepfakes

The advent and widespread accessibility of Generative AI have significantly enhanced the ability to generate vast amounts of content, particularly content that closely mimics the writing style of a specific individual.

For instance, Generative AI can be exploited to create highly convincing phishing emails by leveraging publicly available information and knowledge about a target individual. This approach is increasingly prevalent when tools like Microsoft's Copilot are actively used. If a malicious actor gains access to an individual's email, they can utilize these tools to extract specific writing patterns and craft highly convincing phishing emails with remarkable speed.

Deepfakes also pose a significant threat by facilitating the proliferation of misinformation and challenging our fundamental assumptions about the authenticity of information. It's crucial to acknowledge that these vulnerabilities exploit inherent human tendencies. Human psychology has remained relatively unchanged for millennia, leaving us susceptible to exploitation as these new technologies leverage and twist our ingrained survival mechanisms. While we once considered photographs

to be irrefutable evidence, we now must critically evaluate all information, a task that often proves challenging due to our inherent human limitations.

Emerging Threats

The emergence of Generative AI has introduced new security vulnerabilities. One significant threat is prompt injection, where attackers manipulate prompts to induce the AI model to generate harmful outputs, reveal sensitive information, or perform unintended actions.

A concerning example involves Microsoft's Copilot. Given Copilot's access to user emails (including unopened ones) and potentially other services like internal documentation and calendars, attackers can embed malicious prompts within seemingly innocuous emails. These prompts, often hidden using techniques like white text on a white background, can be interpreted by Copilot, potentially leading to unintended and harmful actions. This vulnerability has already been exploited in real-world scenarios.

Model poisoning is another growing concern. This involves attackers corrupting or altering the training data to influence the model's behaviour subtly. While these initial effects may be subtle, they can have significant and long-term impacts on an organization.

Mitigation

Mitigating these potential vulnerabilities requires a robust security posture. All established security principles and best practices remain highly relevant. Regular security audits and penetration testing are crucial to identify and address any vulnerabilities.

Ethical guidelines and best practices must be established to ensure all AI developments are compliant and explainable. Continuous monitoring is essential to detect threats, anomalies, changes, and abuse within the systems before they escalate. Data access, sharing, and control are critical considerations. Utilizing open-source or on-premise foundational models, rather than relying solely on cloud-hosted commercial models, offers several advantages.

Running AI models on-premises mitigates the potential for data exposure and enables stronger control over model outputs through the use of controlled models and tailored small language models.

The Evolution of Generative AI

Since the initial inflection point of Generative AI becoming viable and accessible, the technology, and crucially, its adoption, has evolved rapidly. A significant recent advancement is the emergence of Generative AI agents. These agents possess the ability to interact with their environment, collect data, learn, and achieve defined goals.

It's crucial to differentiate agents from traditional workflows or automation. Unlike conventional systems, agents can handle unexpected inputs and dynamically determine appropriate responses. Moreover, agents learn and improve over time based on feedback and experience. They operate autonomously, working independently towards predefined goals while also demonstrating adaptability to new objectives and evolving circumstances.

Advanced examples of agents include autonomous vehicles and sophisticated phone answering services. Within the oil and gas industry, agents could be designed to optimize operations by continuously monitoring rigs and plant machinery, analyzing schedules, and proactively taking appropriate remedial and planning actions. These advanced capabilities rely heavily on robust data acquisition and seamless integration. Without a continuous flow of relevant data and the ability to execute actions based on that data, the effective functioning of an agent would be severely limited.

The Reality of Generative AI

Cutting through the hype surrounding Generative AI, it's crucial to acknowledge its limitations. While an incredibly powerful technology, Generative AI is prone to "hallucinations," confidently generating incorrect or misleading information. It can also conflate information and fabricate details.

Furthermore, these models can exhibit limitations when dealing with numerical data and may sometimes struggle with self-consistency. Understanding that Generative AI fundamentally operates as a next-word predictor based on neural networks helps to contextualize these limitations.

While computational costs and speed were initial concerns, ongoing advancements in model efficiency are mitigating these challenges. However, it's essential to acknowledge that in some instances, prioritizing accuracy may necessitate a trade-off in terms of processing speed.

Guardrails and carefully engineered prompts are crucial for mitigating hallucinations and ensuring accurate outputs. These measures often require a trade-off between speed and accuracy, which is acceptable in many applications.

The concept of AI agents represents an exciting evolutionary step. However, many solutions currently labelled as "agents" fall short of true AI agent capabilities. True AI agents, with their ability to interact with their environment, learn, and adapt, are still in their early stages of development.

People First

A key element to any technology adoption is people. No matter how good a technology is, it's useless if people won't use it. Novelty technology interests early adopters but seldom crosses the chasm into mainstream use. The early adopters also drop novelty technology as it fails to add true value beyond the initial interest.

There are key aspects to AI adoption within organisations and they boil down to 4 concepts; understanding, engagement, message and adoption. First, there has to be a clear understanding of what currently happens in a process and why.

Using techniques like contextual inquiries, we can gain valuable insights into how work is done, identifying areas for improvement and potential pitfalls. Engaging with people is crucial for successful adoption. When people feel heard and valued, they're more likely to embrace new technologies.

The language we use can make or break a new technology. By framing it as a tool to enhance existing work, we can increase its chances of adoption. Adoption isn't just about buying new technology. We need to encourage widespread use, even among the most resistant users.

The successful integration of new technologies requires a systematic and iterative approach. By prioritising smaller-scale initiatives, organisations can mitigate risks and foster a culture of innovation. Start small and think big.