Journal of
Petroleum Engineering & Technology
https://journals.stmjournals.com/joPet
ISSN: 2231-1785 (Online)
ISSN: 2321-5178 (Print)
Volume 16, Issue 2, 2026
May–August
DOI (Journal): 10.37591/JoPET

Review JoPET

Artificial Intelligence and Edge Computing in Oil and Gas: Applications, Architectures, and Operational Realities

Mohammed Hazique Shaikh[1, *]

Abstract

Artificial intelligence has arrived in oil and gas, and unlike some previous waves of digital enthusiasm in the sector, this one is sticking. Saudi Aramco analyses approximately 10 billion data point every day and reported USD 4 billion in technology-driven operational gains in 2024. ExxonMobil uses AI to increase shale well output by more than 5 percent. Shell has deployed machine learning across more than 10,000 assets using C3.ai to predict maintenance needs and minimise equipment downtime. These are not pilot projects. They are production deployments at the scale of the world's largest energy companies, and they are changing what it means to operate an upstream or downstream oil and gas facility. This review paper examines the major AI application categories in oil and gas — predictive maintenance, drilling optimisation, reservoir characterisation, pipeline integrity, and process optimisation — and analyses the role of edge computing as the enabling infrastructure that makes real-time AI inference viable in remote, bandwidth-constrained, and operationally sensitive upstream and midstream environments. An original AI-Edge Deployment Matrix (AEDM) is proposed, providing a structured framework for matching AI use cases to the appropriate deployment architecture across the cloud-edge-field continuum. The AI in oil and gas market, valued at USD 5.31 billion in 2024, is projected to nearly triple USD 15.01 billion by 2029.

Keywords: Artificial intelligence, oil and gas, edge computing, predictive maintenance, drilling optimisation, reservoir characterisation, AI-Edge deployment matrix, AEDM, IIoT, digital twin, upstream, downstream

INTRODUCTION

The oil and gas industry has a data problem that most industries would envy. A single offshore platform generates terabytes of sensor data every day — pressure, temperature, flow rate, vibration, chemical composition — from thousands of instrumented points across the facility. A mature oil field might have tens of thousands of producing wells, each generating continuous telemetry. A refinery has hundreds of process units, each with its own sensor network, alarm system, and control loop. The data exists. The problem has always been what to do with it.

For most of the industry's history, the answer has been archive it, look at it occasionally, and react to problems when they become obvious. The maintenance philosophy was largely reactive — fix it when it breaks. The drilling philosophy relied heavily on geological expertise and experience, supplemented by seismic data that took months to process. The process optimisation philosophy used rule-based setpoints established during

***Author for Correspondence**
Mohammed Hazique Shaikh
E-mail: haziqueshaikh@gmail.com

[1]Graduate Student, MS Engineering Management, Northeastern University, Massachusetts, USA
[1]Senior Future Offer Manager, Industrial Automation Division, Schneider Electric, Foxboro, Massachusetts, USA

Received Date: April 30, 2026
Accepted Date: May 22, 2026
Published Date: June 15, 2026

Citation: Mohammed Hazique Shaikh. Artificial Intelligence and Edge Computing in Oil and Gas: Applications, Architectures, and Operational Realities. Journal of Petroleum Engineering & Technology. 2026; 16(2): 1–6p.

commissioning and adjusted infrequently. None of this was negligence. It was the practical reality of working with data volumes that exceeded any analyst's capacity to process manually and processing technologies that were not fast or smart enough to turn raw sensor streams into operational intelligence in real time [1].

AI changes both halves of that equation. Machine learning models can process sensor streams continuously, find patterns invisible to human analysts, and generate actionable predictions at a speed and scale that manual analysis cannot approach. Deloitte reports that edge AI can reduce geo-model creation time from months to hours — a change that directly affects which drilling prospects can be evaluated before market conditions change. The applications are real, the results are documented, and the adoption rate is accelerating, publications related to AI in the oil and gas sector have grown at approximately 15% per year between 2015 and 2024 [2].

What makes this wave different from previous digital transformation initiatives is edge computing. Prior AI deployments in oil and gas were cloud-first: collecting all the data, send it to the cloud, run the models there, receiving the predictions back. This worked for applications tolerant of latency. It broke down for applications that needed real-time inference — drilling parameter optimisation that needs to respond in seconds, pipeline leak detection that needs to trigger an automated isolation valve before a spill reaches a water body, wellhead pressure anomaly detection that needs to alert operators before a well control incident. For these applications, the round-trip to the cloud is too slow and the dependency on internet connectivity is too risky.

Edge computing — deploying compute infrastructure at or near the data source, within the production facility — brings AI inference where it is needed: at the wellhead, at the compressor, at the refinery control room. This review examines the major AI application categories in oil and gas, analyses the edge architectures that enable real-time AI deployment, and proposes the AI-Edge Deployment Matrix (AEDM) as a structured framework for workload placement decisions.

MAJOR AI APPLICATION CATEGORIES IN OIL AND GAS

Predictive Maintenance: The Most Deployed Use Case

Predictive maintenance is where AI has established its strongest operational track record in oil and gas. The business case is straightforward: unplanned equipment failure on an offshore platform or a refinery costs millions per hour in lost production, plus the cost of emergency repair and logistics. If an AI model monitoring vibration signatures on a compressor can predict bearing failure 2 to 3 weeks in advance with reasonable accuracy, it enables planned maintenance during a convenient window, avoids the emergency response, and keeps the facility running [3].

Shell's deployment with C3.ai across more than 10,000 assets is the largest documented example in the sector. The system monitors sensor data continuously, applies machine learning anomaly detection, and generates maintenance recommendations ranked by probability and consequence. BP's AI-based platform reduces the time required for planned maintenance inspections of offshore assets by prioritising which assets need attention based on predicted degradation trajectory, rather than fixed inspection schedules. The financial impact is documented: across the industry, predictive maintenance AI deployments are reducing maintenance costs by 10 to 30 percent and equipment downtime by 15 to 25 percent in the deployments that have published results.

The edge computing requirement for predictive maintenance is moderate: inference needs to complete within a few minutes to hours, not milliseconds. This allows a range of deployment options from cloud-based inference on aggregated data to edge inference on locally processed sensor streams. The choice depends on bandwidth availability, data sovereignty requirements, and the latency tolerance of the maintenance alert workflow [4].

Drilling Optimisation

Drilling is where the financial stakes of AI are highest and where edge computing becomes genuinely critical. A single offshore deepwater well costs USD 50 to 150 million to drill. A drilling optimisation AI that reduces non-productive time by 10 to 15 percent generates direct savings of USD 5 to 20 million per well. BP's deployment of AI for real-time drilling parameter optimisation — adjusting weight on bit, rotary speed, and flow rate based on real-time data analysis — has demonstrably reduced non-productive time across its drilling programme. Deloitte documents that AI can reduce geo-model creation time from months to hours, enabling drilling programmes to incorporate the latest geological intelligence before drilling decisions are finalised.

The edge requirement for drilling optimisation is severe. Geosteering decisions — adjusting the direction of a horizontal wellbore in real time to stay within the target formation — need to be made on timescales of seconds to minutes, based on measurement-while-drilling (MWD) data that is transmitted up the drill string at low bandwidth. Cloud-based geosteering AI cannot respond quickly enough to be operationally useful. Edge compute deployed at the drillsite, processing MWD data locally and running inference in near-real-time, is the required architecture [5].

Reservoir Characterisation and Seismic Interpretation

Seismic data interpretation — the analysis of acoustic wave reflections to build subsurface geological models — generates some of the largest datasets in any industry. A single seismic survey for a large offshore block can produce petabytes of raw data. Traditional processing took months on supercomputers. AI-based seismic interpretation, using deep learning models trained on labelled geological datasets, reduces this processing time to days or hours while improving the accuracy of geological feature identification — particularly for complex structures like salt bodies, fault networks, and thin reservoir intervals.

Chevron's AI-driven seismic interpretation programme has demonstrably improved exploration success rates by reducing the subjectivity inherent in human interpretation of complex seismic volumes. The deployment architecture for this application is predominantly cloud-based: seismic processing is compute-intensive but not time-critical, and the raw data volumes make edge deployment impractical. The cloud is the right architecture here, with the edge serving only as the data acquisition layer [6].

Pipeline Integrity and Leak Detection

Midstream pipeline operators manage thousands of kilometers of infrastructure transporting hydrocarbons under high pressure. Pipeline failures — whether from corrosion, third-party interference, or mechanical fatigue — create environmental, safety, and financial consequences of the first order. AI-based anomaly detection, correlating pressure, flow, and acoustic signals from distributed sensor arrays along the pipeline route, can detect the early signatures of a leak or a structural anomaly hours or days before a failure event.

The edge requirement for pipeline integrity is critical for leak detection: an automated valve isolation system that closes in response to a leak detection alert needs the inference to complete at the field site, not in the cloud, because the time to isolate the failure before a significant spill occurs may be measured in minutes. For longer-horizon integrity monitoring — corrosion rate estimation, remaining life prediction — cloud-based analysis on aggregated historical data is appropriate [7].

Process Optimisation

Process optimisation in refining and petrochemical operations represents one of the highest-value AI application categories in the downstream sector. Modern refineries operate hundreds of interdependent process units — distillation columns, fluid catalytic crackers, hydroprocessing units — each with complex operating variables. Traditionally, process engineers optimise these units using linear programming models updated periodically, with setpoints held constant between optimisation cycles.

AI changes this fundamentally: reinforcement learning and neural network-based optimisation models can adjust operating parameters continuously, responding to feed quality changes, energy price fluctuations, and product demand signals in real time.

ExxonMobil's deployment of AI-driven process optimisation across its refining network delivers percentage-point improvements in yield and energy efficiency — gains that translate directly into hundreds of millions of dollars annually across a large refining portfolio. The edge computing requirement for process optimisation sits in the minutes-to-hours range: fast enough to require local inference at the facility control system level, but not so fast that it demands field-edge deployment. The overriding constraint is data sensitivity: operating conditions and optimisation strategies in competitive refining are commercially confidential, making public cloud deployment unacceptable for most operators and driving adoption of private cloud or on-premises edge architectures.

THE AI-EDGE DEPLOYMENT MATRIX (AEDM)

Different AI applications in oil and gas have different requirements for inference latency, data volume, data sensitivity, and regulatory treatment. The AI-Edge Deployment Matrix (AEDM) provides a structured mapping of use case characteristics to deployment architecture, from fully on-premises edge to fully cloud-hosted inference as shown in Table 1.

Edge AI Infrastructure Requirements for Oil and Gas

Deploying AI inference at the wellsite, pipeline station, or offshore platform introduces infrastructure requirements that are materially different from cloud AI. Compute hardware must be ruggedised for the operating environment — Class I Division 2 (hazardous area) or ATEX Zone 2 certification may be required in hydrocarbon-processing areas. Power availability may be limited: some remote wellsites operate on solar and battery power, with edge computing that must fit within a strict power budget. Connectivity may be satellite-based with high latency and limited bandwidth, making cloud round-trips impractical and edge inference mandatory.

The edge computing hardware that meets these requirements has matured significantly since 2022. ARM-based processors with hardware-accelerated AI inference — NVIDIA Jetson, Intel Movidius, and similar — provide sufficient compute for ONNX-optimised ML models at power levels compatible with industrial edge deployment. Ruggedised edge servers meeting IP65 or IP67 environmental ratings are commercially available from multiple vendors. The software stack — ONNX Runtime for model inference, Docker/Kubernetes for application management, OPC UA for data connectivity — provides the same open, multi-vendor toolchain as data center deployment [8].

Table 1. AI-edge deployment matrix (AEDM) for oil and gas AI applications.

Use Case	Latency requirement	Data sensitivity	Deployment	Rationale
Real-time geosteering	Seconds	High	Field edge	MWD data bandwidth constraints; decision timescale requires local inference
Pipeline leak detection	1–5 minutes	High	Field/pipeline edge	Automated isolation valve needs local trigger; data sovereignty for production rates
Predictive maintenance	Hours	Medium	Facility edge or private cloud	Acceptable latency for maintenance scheduling; sensor data stays on-site
Process optimisation	Minutes	High	Facility edge	Operating conditions competitively sensitive; on-premises inference required
Seismic interpretation	Days	Medium	Public/private cloud	Non-real-time; compute-intensive; anonymisable survey data
Demand forecasting	Hours–days	Low	Public cloud	Market data; no OT sensitivity; full cloud capabilities appropriate

CHALLENGES AND RISKS

Data Quality and Sensor Reliability

The single most consistent operational obstacle to AI deployment in oil and gas is data quality. Machine learning models are only as good as the data they are trained on. In an operating oil field or refinery, sensor data is not clean: sensors drift, fail, saturate, and report physically impossible values. Transmitters get knocked offline by maintenance activities. Communication links drop and create gaps. A predictive maintenance model trained on clean historical data from a laboratory environment and deployed on noisy production sensor streams will generate false positives at rates that operators cannot tolerate, causing the model to be disabled — not because the AI is wrong, but because the data quality in the field was never validated as a prerequisite for deployment [9].

Organisational Change and Workforce Capability

The AI tools are advancing faster than the organisations deploying them. A drilling engineer who has spent fifteen years developing geological intuition about a specific basin does not automatically trust an AI model that contradicts that intuition, even when the model is statistically correct. Shell and BP both report that the most significant challenge in deploying AI-based predictive maintenance is not the model — it is convincing experienced maintenance engineers that the model output is trustworthy enough to act on without independent verification. Explainability mechanisms — SHAP, LIME, attention visualisation — help, but they require investment in the model design and the operator interface that is often cut from deployment budgets [10].

FUTURE DIRECTIONS

Three technology trajectories are expected to reshape AI deployment in oil and gas over the next five years. The first is the maturation of large language models (LLMs) for industrial applications. Operators are beginning to evaluate LLMs for accelerating engineering document retrieval, procedure compliance checking, and real-time operator decision support. Models capable of ingesting maintenance histories, process datasheets, and operational logs to provide contextually relevant guidance to field operators represent a significant advance over current rule-based expert systems.

The second trajectory is the convergence of digital twin technology with edge AI. A digital twin — a high-fidelity simulation model of a physical asset, updated in real time with sensor data — provides the contextual model that makes AI predictions interpretable and actionable. When a predictive maintenance alert fires, a digital twin can show the operator not just that an anomaly is present, but which component is degrading, what the failure mechanism is, and how much remaining useful life is projected under different operating scenarios. Integrating this capability at the facility edge, rather than relying on cloud roundtrips for simulation updates, is a near-term roadmap priority for major operators.

The third trajectory is the development of AI governance and assurance frameworks specific to safety-critical process applications. As AI models take on roles adjacent to safety instrumented functions — leak detection, emergency shutdown triggering, well control decision support — regulatory frameworks governing their qualification, validation, and change management are evolving. The International Society of Automation (ISA) and the International Electrotechnical Commission (IEC) are developing standards defining how AI systems must be documented, tested, and audited in process safety contexts. Early engagement with these frameworks is becoming a prerequisite for operators deploying AI in high-consequence applications. Collectively, these directions point toward AI systems that are not only more capable but more transparent, governable, and safe.

CONCLUSION

AI in oil and gas is no longer a future story. It is a present operational reality at the world's largest energy companies, and the documented financial impacts — billions in operational gains for Saudi Aramco, percentage-point improvements in production yields for ExxonMobil, measurable reductions in unplanned maintenance costs for Shell — establish the business case beyond reasonable doubt. The market is set to triple from USD 5.31 billion to USD 15.01 billion by 2029.

What the AEDM makes clear is that AI deployment in oil and gas is not a single architectural decision but a portfolio of decisions, each matching a use case to the infrastructure that best serves its latency, data sensitivity, and regulatory requirements. Seismic interpretation belongs in the cloud. Real-time geosteering belongs at the field edge. Predictive maintenance sits in the middle, where a hybrid edge-private-cloud architecture — local inference for time-sensitive alerts, cloud aggregation for long-horizon models — provides the best of both architectures. Getting those placement decisions right is as important as getting the AI models right. Both require deliberate engineering.

DECLARATION OF INTEREST

The author declares that there is no conflict of interest regarding the publication of this manuscript. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Acknowledgement

The author acknowledges the broader industrial automation and digital infrastructure research community whose published work informs this review. Practitioner experience in industrial automation product management provided the applied context for the analysis presented.

REFERENCES

1. Anifowose FA. Artificial intelligence application in reservoir characterization and modeling: whitening the black box. InSPE Saudi Arabia section Young Professionals Technical Symposium 2011 Mar 14 (pp. SPE-155413). SPE.
2. Wang H. Sustaining the AI Revolution With Analog Compute-in-Memory [Industrial and Governmental Activities]. IEEE Computational Intelligence Magazine. 2026 Jan 14;21(1):11-2.
3. Póvoas MD, Moreira JF, Neto SV, Carvalho CA, Cezario BS, Guedes AL, Lima GB. Artificial intelligence in the oil and gas industry: Applications, challenges, and future directions. Applied Sciences. 2025 Jul 16;15(14):7918.
4. George AS, George AH, Baskar T. Edge computing and the future of cloud computing: A survey of industry perspectives and predictions. Partners Universal International Research Journal. 2023 Jun 20;2(2):19-44.
5. Garcia-Amate A, Wesley D, Puffer SM. O&G 4.0: Leveraging Technology for Better Environmental Outcomes. InEfficiency and Optimization in Business Practice: Enhancing Productivity, Innovation and Cost-Efficiency for Sustainable Growth 2025 Sep 6 (pp. 21-42). Cham: Springer Nature Switzerland.
6. Mardanov E, Mavlutova I, Sloka B. Financial Performance Outcomes of AI-Adoption in Oil and Gas: The Mediating Role of Operational Efficiency. Journal of Risk and Financial Management. 2026 Jan 6;19(1):44.
7. Hussain M, Alamri A, Zhang T, Jamil I. Application of artificial intelligence in the oil and gas industry. InEngineering applications of artificial intelligence 2024 Feb 20 (pp. 341-373). Cham: Springer Nature Switzerland.
8. Knebel FP, Trevisan R, do Nascimento GS, Abel M, Wickboldt JA. A study on cloud and edge computing for the implementation of digital twins in the Oil & Gas industries. Computers & Industrial Engineering. 2023 Aug 1;182:109363.
9. Zakizadeh M, Zand M. Revolutionizing Oil & Gas: A Comprehensive Review of Smartening Technologies in the oil & Gas Industry. InProceedings of the International Conference of New Technologies in Oil, Gas and Petrochemical Engineering in Iran 2024 Oct (pp. 22-23).
10. Póvoas MD, Moreira JF, Neto SV, Carvalho CA, Cezario BS, Guedes AL, Lima GB. Artificial intelligence in the oil and gas industry: Applications, challenges, and future directions. Applied Sciences. 2025 Jul 16;15(14):7918.

