

THE APPLICATION OF AUTONOMOUS AI AGENTS IN UPSTREAM OIL AND GAS SUPPLY CHAIN DECISION-MAKING

Ismail Al-Zadjali

Muscat University, Oman

240133@muscatuniversity.edu.om

ABSTRACT

The upstream oil and gas industry is one of the most complex supply chain sectors in the world, with multi-tier vendor networks, volatile commodity prices, logistically challenging field operations and strict safety and regulatory standards. Traditional methods of decision-making, including enterprise resource planning (ERP) systems and traditional analytics, are increasingly unable to deliver timely action and predictive intelligence that today's upstream operations require. They are an embodiment of a new paradigm in the conception and operationalization of upstream supply chain decisions, autonomous artificial intelligence (AI) agents, goal-directed computational entities that can perceive environmental states, reason over multi-dimensional data, formulate adaptive plans and execute actions without needing to be continuously monitored by humans. In this paper, the author shares insight on the use of autonomous AI agents for the key business decisions in the upstream supply chain, ranging from exploration data management to predictive asset maintenance, including drilling operations planning, procurement and vendor selection, inventory and materials management, and logistics optimization. The paper builds upon a systematic literature review of existing literature, industry case studies of large-scale operators such as Saudi Aramco, Shell, ExxonMobil, ConocoPhillips and ADNOC, and an analysis of emerging multi-agent system (MAS) and large language model (LLM) agent architectures to develop a detailed taxonomy of agent types and their capabilities to specific upstream supply chain functions. Two tables summarizing the agents typology and the implementation outcomes of the industry are provided as illustrative examples, as well as conceptual framework diagrams. Ongoing challenges in the field of data heterogeneity, cybersecurity exposure, regulatory compliance, explainability deficits and human-agent trust are also examined in the paper, followed by a strategic implementation roadmap. The results validate the emergence of autonomous AI agents from experimental prototypes to become part of an organization's infrastructure and the benefits for organizations that are systematically integrating agentic intelligence into their supply chain architecture.

Keywords:

Autonomous AI Agents, Upstream Oil and Gas, Supply Chain Decision-Making, Multi-Agent Systems, Reinforcement Learning, Large Language Models, Digital Transformation, Predictive Analytics, Procurement Optimisation, Drilling Intelligence.

1. INTRODUCTION

The upstream oil and gas sector covers all the activities from geophysical exploration and seismic surveying to reservoir appraisal, field development planning, drilling, and primary production. It is one of the most capital-intensive, technically challenging, and risky environments in the global economy. The supply chains that enable upstream operations are fundamentally different from manufacturing or retail, across multiple continents and ocean basins; include several thousand unique stock-keeping units (SKUs); have multiple tiers of contractors and vendors working under very special terms and conditions; and are susceptible to sudden and unexpected disruptions from geopolitical events, natural phenomena, equipment failures and commodity price volatility [1]. The choices that have to be made to run these supply chains involve different time horizons – from multi-year decision-making about vendor qualification and contract structuring, to weeks and months for last-minute changes to orders and materials, to real-time calculations that take mere minutes or hours.

Traditional supply chain management models, even with enterprise resource planning (ERP) systems, statistical demand forecasting, and business intelligence dashboards, are inherently constrained in dealing with this complexity. They are primarily reactive, on a periodic refresh cycle rather than continuous, and consume a lot of human bandwidth, which is expensive and prone to cognitive biases and bandwidth limitations. With the advent of Artificial Intelligence (AI) and, in particular, the new class of AI systems called autonomous agents, this new approach is qualitatively different. An autonomous AI agent is a computational agent that senses its environment

and acquires knowledge about the current state of the world, thinks about the world state given those perceptions to develop goals and plans, and takes actions, either through physical actuators or software interfaces, to achieve those goals and plans, in an autonomous manner to reach autonomy in comparison to a conventional algorithmic automation [8]. While rule-based expert systems operate according to a pre-programmed logic tree, modern AI agents are built on machine learning, deep reinforcement learning, and, most importantly, the powerful natural-language reasoning capabilities of large language models (LLMs). They can work with a new decision context, one that is inherently ambiguous and evolves over time [3,4].

The past five years have witnessed a pronounced acceleration in the deployment of AI agents in upstream operations. Landmark implementations—including Saudi Aramco's enterprise-scale LLM agent processing ten billion data points daily, ADNOC's ENERGYai autonomous gas-lift control loops, Baker Hughes' Leucipa multi-agent production optimization platform, and Shell's AI-driven seismic processing that delivered a 99% reduction in processing time demonstrate that agentic AI is no longer confined to research laboratories but is actively reshaping the operational economics of upstream production [6,7]. These implementations span multiple supply chain functions: exploration data management, drilling optimization, predictive maintenance, logistics coordination, and procurement decision support.

Nevertheless, significant knowledge gaps persist in the academic and practitioner literature. Most existing reviews address AI in oil and gas at a general or function-specific level, without systematically mapping the full spectrum of autonomous agent architectures to the breadth of upstream supply chain decision domains, or providing a structured analysis of implementation challenges and strategic enablers. This paper addresses those gaps through four core contributions: (i) a comprehensive taxonomy of autonomous AI agent types and their technical characteristics as applied to upstream supply chain contexts; (ii) a systematic mapping of agent capabilities to upstream supply chain decision functions; (iii) a synthesis of documented industry case studies evidencing operational outcomes; and (iv) a structured analysis of implementation challenges and a proposed strategic roadmap for scaling agentic AI across upstream supply chain organizations.

The remainder of the paper is organized as follows. Section 2 provides a structured literature review of AI agent theory, supply chain management in the oil and gas industry, and existing empirical evidence. Section 3 presents the agent typology and classification framework. Section 4 maps agent applications to upstream supply chain decision domains. Section 5 analyses industry case studies and presents the comparative evidence tables. Section 6 discusses implementation challenges and risk factors. Section 7 proposes a strategic implementation roadmap. Section 8 presents a discussion of future research directions, and Section 9 concludes the paper.

2. LITERATURE REVIEW

2.1 The Theory and Architecture of Autonomous AI Agents

The theoretical foundations of autonomous AI agents trace to classical artificial intelligence research, particularly the Beliefs-Desires-Intentions (BDI) model of rational agency developed through the 1980s and formalized by Wooldridge [9] and Russell and Norvig [8]. In the BDI framework, an agent maintains a set of beliefs about its environment (derived from perception), a set of desires or goals it seeks to achieve, and a set of intentions selected plans of action it is currently committed to executing. This tripartite architecture provides a principled basis for designing agents capable of deliberative reasoning and adaptive goal pursuit in dynamic environments.

Contemporary AI agent architectures extend the classical BDI model in significant ways. The advent of deep reinforcement learning (RL), formalized by Sutton and Barto [10] and demonstrated at superhuman performance levels by Mnih et al. [11] and Silver et al. [12], has enabled agents to learn optimal decision policies through iterative environmental interaction, without requiring complete a priority world models. Reinforcement learning agents receive scalar reward signals contingent on the outcomes of their actions and update their value function or policy representations accordingly a mechanism highly suited to supply chain environments where objective functions (cost minimization, delivery time reduction, utilization maximization) can be quantitatively specified.

Multi-agent systems (MAS), in which networks of interacting agents collectively address problems too complex for any single agent, represent a further theoretical advance [9,13]. The seminal work of Lowe et al. [13] on multi-agent actor-critic frameworks and Rashid et al. [14,38] on QMIX a monotonic value function factorization method enabling cooperative multi-agent learning established key algorithmic foundations. In supply chain contexts, MAS architectures are particularly relevant because supply chains are inherently distributed systems involving multiple heterogeneous decision-making entities (procurement offices, logistics coordinators, vendor representatives, field operations teams) whose decisions are coupled through shared resources and information flows [25,45].

The most recent wave of agent development is centred on large language model (LLM)-based agents [3,4,5]. LLMs such as GPT-4 and LLaMA [37] possess emergent capabilities for instruction following, multi-step

reasoning, tool use, and multi-turn dialogue that render them powerful cognitive cores for autonomous agents. The ReAct framework [17] demonstrated that LLMs could interleave reasoning traces and tool-execution actions in a coherent loop, enabling the completion of complex tasks across heterogeneous information sources. Park et al. [18] extended this to multi-agent simulations of human social behaviour, while Liu et al. [33] developed AgentBench to systematically evaluate LLM agent capabilities across standardized task environments. Song et al. [32] specifically investigated pre-trained LLMs adapted for industrial control applications, establishing the relevance of this paradigm to operational engineering domains.

2.2 Supply Chain Management in Upstream Oil and Gas

The upstream oil and gas supply chain has been described in the literature as a system of extreme complexity, characterized by high asset specificity, geographically distributed demand nodes, stochastic lead times, and heterogeneous regulatory environments [1,2]. The supply chain encompasses four principal phases: exploration and appraisal (geological surveying, seismic acquisition, data processing, well appraisal); drilling and completions (rig procurement, casing and tubular management, directional drilling services, completion operations); production operations (artificial lift, chemical injection, production chemistry, separator management); and decommissioning. Each phase has distinct procurement patterns, lead time profiles, and criticality structures that shape supply chain risk [40].

Earlier reviews of AI in oil and gas supply chain contexts noted the predominance of statistical and rule-based approaches and highlighted the emerging role of machine learning [1,2]. Helo and Hao [19] examined AI in operations management more broadly, identifying demand forecasting, process automation, and anomaly detection as the functions most amenable to AI enhancement. Koroteev and Tekic [1] observed that early AI planning applications using optimization and machine learning methods executed by autonomous robots and intelligent agents were emerging in the oil and gas industry but remained nascent. Subsequent research by Patidar et al. [36] and Delpisheh et al. [26] documented the integrating role of IoT sensor networks and digital twin technologies in providing the real-time data substrates necessary for AI agent operation.

Supply chain resilience has emerged as a critical theme in the post-COVID-19 literature. Ivanov and Dolgui [21] and Dolgui et al. [22] developed theoretical frameworks for understanding ripple effects and disruption propagation through complex supply networks frameworks directly applicable to the geopolitically exposed upstream oil and gas context. Smyth et al. [23] conducted a systematic literature review of AI and prescriptive analytics for supply chain resilience, identifying autonomous decision support as a key emerging capability. Khan and Islam [27] examined the specific application of digital twins which function as the environmental perception and simulation substrate for AI agents in oil and gas supply chains, documenting their capacity to compress scenario-planning cycles from weeks to hours.

2.3 AI Agent Applications in Upstream Operations: Empirical Evidence

Empirical evidence of AI agent applications in upstream operations has accumulated rapidly since 2020. In drilling optimization, Erge and van Oort [28] demonstrated the effectiveness of expert-guided reinforcement learning for rate-of-penetration (ROP) optimization across multiple bit runs and rigs, achieving generalization across heterogeneous geological formations a key capability for agents deployed across diverse field locations. Hegde and Grey [29] documented the use of machine learning to improve drilling efficiency in adjacent wells, establishing the data inheritance patterns that agentic systems can exploit. Osarogiagbon et al. [30] reviewed machine learning for hazardous event prediction in drilling, encompassing well-control events, stuck pipe, and blowout precursor detection functions where autonomous alert-and-response agents can measurably reduce safety incidents.

In production optimization, Chen et al. [39] documented RL agent applications for gas-lift optimization in offshore operations, demonstrating the capacity of continuous learning agents to outperform both human expert rules and conventional optimization solvers in dynamic multi-well production environments. Park et al. [48] developed hybrid deep learning models for predictive maintenance of offshore production equipment, establishing the technical foundations for autonomous maintenance scheduling agents. Agrawal [35] examined generative AI for supply chain management, identifying use cases in demand sensing, exception management, and supplier communication automation that are directly transferable to upstream procurement contexts.

At the industry deployment level, publicly documented implementations by Saudi Aramco, Shell, ADNOC, ExxonMobil, ConocoPhillips, and Baker Hughes [6,7] collectively constitute a substantial evidence base. These deployments are reviewed and synthesized in Section 5 of this paper. Smyth et al. [23] and Zhu et al. [24] provide additional evidence that AI and multi-agent systems can deliver measurable improvements in supply chain performance in adjacent industrial contexts. Ning and Xie [25] survey multi-agent reinforcement learning for

resource allocation optimization, establishing the algorithmic repertoire available for application to upstream materials and logistics management.

3. TAXONOMY OF AUTONOMOUS AI AGENTS IN UPSTREAM SUPPLY CHAIN CONTEXTS

3.1 Classification Dimensions

A structured classification of AI agents applicable to upstream supply chain decision-making requires analysis along several orthogonal dimensions: (i) the agent's architectural paradigm (reactive, deliberative, hybrid, learning, or multi-agent); (ii) the decision-making horizon targeted (operational/real-time, tactical/medium-term, or strategic/long-term); (iii) the degree of autonomy exercised (human-in-the-loop, human-on-the-loop, or fully autonomous); and (iv) the primary data modality processed (structured sensor data, tabular enterprise data, unstructured text, or multi-modal combinations). Table 1 below presents a comparative typology of the principal agent classes applicable to upstream supply chain contexts, summarising their core functionality, representative upstream applications, autonomy levels, and primary limitations [8,9,25,45].

Table 1: Taxonomy and Comparative Analysis of Autonomous AI Agent Types in Upstream Oil and Gas Supply Chain Applications

Agent Type	Core Functionality	Upstream Application	Decision Autonomy Level	Key Limitation
Reactive Agent	Responds directly to real-time sensor stimuli	Wellhead pressure & flow monitoring	Low – rule-based	Cannot plan ahead
Deliberative Agent	Uses internal world models to reason and plan	Drilling trajectory optimization	Medium – model-driven	High computational demand
Hybrid Agent	Combines reactive and deliberative modules	Integrated HSE & production scheduling	Medium–High	Complex integration overhead
Learning Agent (RL)	Updates policy through environmental feedback	Dynamic procurement bidding & routing	High – adaptive	Data-hungry; slow convergence
Multi-Agent System (MAS)	Networked agents collaborate/compete	Cross-functional supply chain coordination	Very High – distributed	Coordination complexity
LLM-Based Agent	Natural-language reasoning over unstructured data	Contract analysis, vendor onboarding	Medium – human-validated	Hallucination risk

3.2 Reactive Agents

Reactive agents constitute the most elemental tier of the agent hierarchy. Operating without internal world models or planning mechanisms, reactive agents implement direct stimulus-response mappings: when a specific environmental condition is detected (e.g., wellhead pressure exceeding a defined threshold), the agent executes a predefined response (e.g., triggering an alert, initiating a valve closure, logging the event to a maintenance system). In upstream supply chain contexts, reactive agents are extensively deployed for real-time monitoring SCADA-integrated agents that monitor pipeline integrity, chemical injection rates, and production separator performance [1,40]. Their principal advantages are computational efficiency and deterministic predictability; their limitation is the inability to anticipate future states or adapt to novel conditions outside the programmed response space.

3.3 Deliberative and Hybrid Agents

Deliberative agents maintain an explicit internal model of the world a symbolic or numeric representation of relevant environmental states, causal relationships, and objective functions and employ search or optimization

algorithms to derive plans for achieving specified goals [8]. In upstream supply chain contexts, deliberative agents find application in drilling trajectory optimization (where the agent reasons over geological models, wellbore trajectory constraints, and equipment specifications to generate multi-segment well plans), logistics route planning (where the agent optimizes multi-stop delivery sequences subject to load constraints, hazardous materials regulations, and road or maritime access restrictions), and contract portfolio analysis (where the agent evaluates vendor bids against multi-criteria weighting schemes).

Hybrid agents integrate reactive and deliberative modules within a unified architecture, typically organized in a layered structure in which reactive behaviours handle urgent immediate responses. In contrast, deliberative processes manage medium- and long-term planning [9]. The layered BDI architecture in which reactive behaviours at the perception layer are supervised by deliberative planning processes at a higher cognitive layer has been shown to deliver robust, practical performance in complex real-world environments. In upstream contexts, hybrid agents are deployed in integrated HSE monitoring and production scheduling applications, where the system must simultaneously react to real-time safety signals and maintain long-horizon production plans.

3.4 Learning Agents and Reinforcement Learning

Learning agents extend the deliberative paradigm by incorporating mechanisms for updating their internal models and decision policies based on experienced outcomes [10,11]. Reinforcement learning agents, in particular, have demonstrated remarkable capability in upstream optimization tasks. An RL agent deployed for rate-of-penetration optimization, for instance, receives state observations (weight-on-bit, rotary speed, mud flow rate, formation hardness indicators) and selects drilling parameter adjustments to maximize a reward signal calibrated to penetration rate, bit wear, and wellbore quality metrics [28]. Through thousands of simulated or real interaction cycles, the agent learns to generalize effective drilling strategies across diverse geological formations a capability that human drillers, constrained by finite experience bases, cannot replicate at scale [29].

The application of RL to upstream procurement introduces distinctive dynamics. A procurement agent can be trained to optimize bidding strategies across multiple vendor relationships, balancing cost minimization against delivery reliability, quality assurance scores, and supply continuity risk. By modelling vendor behaviour as a stochastic environment and the procurement outcome as a multi-step Markov decision process, the RL agent learns procurement policies that outperform both rule-based approaches and human expert judgment in complex, time-pressured sourcing environments [35,50].

3.5 Multi-Agent Systems and Distributed Decision-Making

Multi-agent systems (MAS) represent the most powerful and architecturally complex tier of the agent taxonomy, deploying networks of interacting agents that collectively address supply chain decision problems too large or structurally complex for single-agent treatment [9,13,45]. In upstream oil and gas supply chains, MAS architectures are particularly well-suited to the inherent distributed structure of the decision environment: exploration data management, drilling operations, procurement, logistics, and maintenance represent functionally distinct sub-systems, each with its own data environment, decision cadence, and stakeholder community, yet each coupled to the others through resource flows and information dependencies [2,25].

A representative MAS architecture for upstream supply chain optimization might deploy: (i) exploration agents that process seismic data and generate prospect portfolios; (ii) procurement agents that solicit and evaluate vendor bids for materials and services required to develop selected prospects; (iii) logistics agents that plan and monitor the multi-modal transport of materials and equipment to field locations; (iv) drilling operations agents that optimize in-situ drilling parameters in real time; and (v) a supervisory coordination agent that monitors the collective performance of the system against field development objectives and reallocates resources across sub-agents when priorities shift. Algorithms such as QMIX [14] and multi-agent actor-critic frameworks [13] provide the mathematical foundations for training such systems to coordinate effectively.

3.6 LLM-Based Agents

The emergence of large language model-based agents represents the most recent and rapidly evolving frontier of autonomous AI agent development [3,4,5,16,17]. LLM agents leverage the language understanding, multi-step reasoning, instruction-following, and tool-use capabilities of foundation models to address upstream supply chain tasks involving unstructured information. This category encompasses the vast majority of upstream procurement documentation, regulatory filings, incident reports, vendor communications, and technical specifications.

In upstream supply chain contexts, LLM agents are being applied to contract analysis and exception identification (automatically parsing multi-hundred-page frame agreements against compliance checklists and flagging deviations), vendor onboarding automation (extracting and validating qualification data from heterogeneous vendor submission formats), incident report analysis (clustering and thematically analyzing large corpora of HSE incident reports to surface systemic risk patterns), and materials management (translating technical requisitions

expressed in engineering language into structured purchase order specifications). The ReAct framework [17] enables LLM agents to interleave reasoning and action in a closed loop for instance, querying an ERP system for inventory levels, reasoning about whether current stock covers projected demand given lead time estimates, and issuing a purchase recommendation all within a single autonomous task execution cycle.

4. MAPPING AI AGENT CAPABILITIES TO UPSTREAM SUPPLY CHAIN DECISION DOMAINS

4.1 Exploration and Seismic Data Management

The exploration phase of the upstream supply chain generates extraordinary volumes of heterogeneous data, including two- and three-dimensional seismic surveys, well logs, mud logs, core sample analyses, satellite imagery, and regional geological maps. Managing, integrating, and interpreting this data has historically required large teams of geoscientists and significant processing time bottlenecks that delay prospect evaluation and increase exploration costs [1,26]. Autonomous AI agents are transforming this landscape by automating the perception, classification, and preliminary interpretation of seismic volumes. Shell's documented deployment of AI-powered seismic processing agents achieved a 99% reduction in processing time [6], compressing decision cycles that previously required months into days and enabling exploration teams to evaluate far larger prospect portfolios within fixed budgets.

From a supply chain perspective, faster and more accurate exploration data management directly improves the quality of forward supply chain planning: wells approved for drilling earlier generate earlier material requirements, enabling longer procurement lead times, better vendor selection, and reduced premium freight costs. Autonomous seismic agents thus function as demand signal generators for the broader upstream supply chain, and their accuracy improvements propagate as quality improvements throughout all downstream supply chain decisions.

4.2 Drilling Operations Planning and Real-Time Optimisation

Drilling represents the highest per-unit-time cost activity in upstream operations, and the supply chain function is most sensitive to the quality of real-time decisions. Drilling supply chains must synchronize the delivery of drill bits, casing strings, completion equipment, chemicals, and specialized services each with distinct lead times and criticality profiles against a drilling schedule that is itself subject to real-time revision based on subsurface conditions encountered [28,29]. AI agents operating in the drilling domain span the full spectrum from reactive (real-time wellbore instability alert agents) through deliberative (daily drilling programme planning agents) to reinforcement learning-based (continuous rate-of-penetration optimization agents).

Erge and van Oort [28] demonstrated that a generalized RL agent for ROP optimization, trained across multiple bit runs and rig types, could outperform both human drillers and conventional model-based optimizers by adapting its parameter recommendations in real time to formation variability. This capability has direct supply chain implications: by extending bit life and reducing unnecessary wiper trips and remediation operations, the RL drilling agent reduces emergency consumable procurement events the highest-cost, shortest lead-time category of upstream procurement. ConocoPhillips' deployment of RL-based ROP optimization resulted in significant drilling cost reductions that cascade directly into improved supply chain economics [6].

4.3 Procurement and Vendor Selection

Procurement in upstream oil and gas is structurally complex: the category landscape encompasses thousands of unique materials and services, from commodity low-value items to highly engineered, single-source components; vendor qualification requirements are stringent and jurisdiction-specific; and procurement decisions must balance cost, lead time, quality, and supplier financial stability across volatile market conditions [7,19,35]. Autonomous AI agents are being applied across the procurement value chain: specification-to-purchase-order matching agents that translate engineering material requisitions into structured procurement specifications; vendor scoring and ranking agents that maintain dynamic vendor performance records and recommend optimal sources for each category; contract compliance monitoring agents that track purchase order progress against contractual commitments and alert procurement managers to emerging delivery risks; and market intelligence agents that monitor commodity price indices, freight cost benchmarks, and geopolitical risk indicators to inform strategic sourcing decisions.

The application of RL agents to procurement bidding dynamics introduces a particularly powerful capability: the ability to learn and exploit patterns in vendor bidding behaviour across repeated procurement interactions. An RL procurement agent that has observed hundreds of historical bidding events can model each vendor's likely bid distribution and timing behaviour, enabling the buyer to sequence information releases and negotiation touchpoints in ways that systematically extract better commercial outcomes without violating procurement ethics or regulatory requirements [35,42]. LLM-based agents further enhance procurement intelligence by enabling the

automated extraction of compliance and capability information from heterogeneous vendor qualification documentation a task that conventionally requires significant analyst time [4,5].

4.4 Inventory and Materials Management

Inventory management in upstream oil and gas is characterized by the simultaneous presence of extremely high carrying costs (for large, capital-intensive items held in remote field warehouses) and extremely high stockout costs (for critical spare parts whose unavailability can shut down producing wells at daily rates measured in tens or hundreds of thousands of dollars) [40]. The heterogeneity of demand patterns with most items exhibiting intermittent, lumpy demand driven by maintenance events renders conventional statistical forecasting methods inadequate and creates persistent over- or under-stocking patterns.

AI agents operating in the inventory domain typically implement a hierarchy of functions: demand sensing agents that integrate real-time asset condition data (from IoT sensors and condition monitoring systems) with historical consumption records to generate probabilistic demand forecasts; stock-level monitoring agents that continuously assess inventory positions against dynamic safety stock targets calibrated to real-time lead time and demand uncertainty estimates; and replenishment decision agents that generate purchase recommendations and, in fully autonomous implementations, initiate purchase orders within pre-approved authority limits. Patidar et al. [36] document the integration of IoT, digital twin, and AI agent technologies in upstream inventory management, demonstrating that this integrated approach can reduce inventory carrying costs by 15–25% while simultaneously reducing critical stockout events.

4.5 Logistics and Transportation Optimisation

Upstream logistics encompasses the coordination of materials and equipment movements across highly complex multi-modal networks: road freight to remote onshore well sites, marine vessel operations to offshore platforms, helicopter logistics for personnel and urgent cargoes, and specialized heavy haulage for large capital equipment [1,40]. The combinatorial complexity of optimizing these movements subject to vessel and helicopter scheduling constraints, port and heliport slot availability, weather window restrictions, hazardous materials regulations, and customs requirements exceeds the capacity of human planners operating with conventional tools.

Multi-agent logistics optimization systems decompose this complexity across specialized agents: route planning agents that optimize freight sequences and consolidation across vessels or road transport cycles; vessel scheduling agents that match cargo requirements against fleet availability while minimizing non-productive time and fuel consumption; customs and documentation compliance agents that prepare and validate shipping documentation against destination country regulatory requirements; and disruption response agents that detect weather window closures, vessel breakdowns, or road access events and replan logistics sequences in real time to minimize supply continuity impacts. A leading pipeline operator's deployment of a network of leak-detection and autonomous rerouting agents monitoring a 1,200-mile pipeline continuously and autonomously rerouting flow upon anomaly detection illustrates the maturity of multi-agent logistics intelligence in operational energy infrastructure [12_xenon].

4.6 Predictive Maintenance and Asset Integrity

Predictive maintenance represents perhaps the most commercially proven application of AI agents in upstream operations. The value proposition is straightforward: by detecting equipment degradation signatures from continuous sensor data streams, AI agents can schedule maintenance interventions before failure events occur, eliminating or reducing unplanned production downtime, which in offshore contexts can carry daily revenue loss exposure exceeding USD 1 million per producing well [48]. Machine learning agents trained on historical sensor data and failure event records learn to recognize the multivariate signatures of impending failures in rotating equipment (pumps, compressors, turbines), wellbore control equipment (valves, chokes), and production treatment systems [19,34].

The supply chain implications of autonomous predictive maintenance agents extend beyond the immediate maintenance event. By generating advance notice of likely failure events with probabilistic time-to-failure estimates, predictive agents enable planned procurement of maintenance parts and mobilization of specialist service contractors transforming emergency, premium-freight procurement events into planned, competitive procurement exercises. Park et al. [48] documented hybrid deep learning models for predictive maintenance of offshore equipment that demonstrated significant lead-time advantages over reactive maintenance practices. The integration of predictive maintenance agent outputs with inventory management agents creates a closed-loop demand signal system for maintenance materials and services, substantially reducing both inventory carrying costs and emergency procurement expenditures.

5. INDUSTRY CASE STUDIES AND COMPARATIVE EVIDENCE

5.1 Overview of Documented Industry Deployments

The following analysis synthesizes publicly documented deployments of autonomous AI agents by major upstream oil and gas operators, drawing on industry publications, corporate technical reports, and peer-reviewed case studies. Table 2 presents a comparative summary of seven landmark deployments, spanning exploration data management, production optimization, supply chain coordination, and pipeline operations. The cases collectively illustrate the progression from isolated function-specific AI tools to integrated, agentic systems operating across multiple supply chain domains.

Table 2: Documented Industry Deployments of Autonomous AI Agents in Upstream Oil and Gas Supply Chain Functions

Company / Platform	AI Agent Application	Supply Chain Function	Reported Outcome	Reference Year
Shell (ADNOC Cooperation)	Autonomous seismic data processing agents	Exploration data management	99% reduction in seismic processing time	2024
Saudi Aramco	Enterprise LLM agent for daily analytics	Data-driven decision-making	USD 4 billion tech-driven gains (2024)	2024
ExxonMobil	ML agent for shale well optimization	Production engineering	>5% increase in shale well output	2023
ConocoPhillips	RL agent for rate-of-penetration optimization	Drilling operations	Significant drilling cost reduction	2023
ADNOC (ENERGYai)	Agentic workflow for gas-lift control	Production optimization	Autonomous gas-lift control loops established	2024
Baker Hughes (Leucipa)	Multi-agent production optimization platform	Upstream production scheduling	Continuous optimization cycle enabled	2023
Pipeline Operator (Anonymous)	Leak-detection & rerouting agent network	Midstream pipeline integrity	Auto-rerouting on anomaly detection; MTTR reduced	2025

5.2 Saudi Aramco: Enterprise-Scale LLM Agent for Data-Driven Decision-Making

Saudi Aramco's deployment of an enterprise LLM agent processing approximately ten billion data points per day [6,13_blackridge] represents the most extensive documented application of autonomous AI to upstream data management and decision support. The system integrates structured production data, well performance records, supply chain transactional records, and unstructured operational communications into a unified agent environment, enabling the generation of natural-language decision recommendations across multiple operational domains simultaneously. Reported technology-driven gains of USD 4 billion in 2024 [7] span production optimization, maintenance scheduling, and supply chain efficiency improvements confirming the commercial materiality of enterprise-scale AI agent deployment.

The Aramco implementation illustrates a critical design principle for LLM-based supply chain agents: the importance of domain-grounded reasoning. Rather than deploying a general-purpose LLM without operational context, Aramco's system integrates the LLM cognitive core with real-time data feeds, domain-specific knowledge bases, and structured output validators. This configuration constrains the agent's reasoning within operationally validated boundaries and reduces the risk of hallucination [5,32]. This architecture exemplifies the human-AI collaborative model in which AI agents function as strategic advisors, generating recommendations that human experts validate before execution.

5.3 ADNOC ENERGYai and Baker Hughes Leucipa: Agentic Workflow Platforms

ADNOC's ENERGYai platform and Baker Hughes' Leucipa production optimization system represent the leading examples of purpose-built agentic workflow architectures for upstream supply chain and production management [6]. Both systems implement the perception-reasoning-action loop continuously ingesting sensor data from

producing wells, reasoning over multi-well production system interactions using reservoir and network models, and autonomously adjusting gas-lift injection rates, choke settings, and production routing configurations to maximize field-level output within constraint envelopes.

From a supply chain perspective, the significance of these deployments lies in their demonstration that closed-loop autonomous optimization can operate at field scale across dozens of producing wells simultaneously a task that would require a team of production engineers operating 24 hours per day if performed manually. The continuous optimization cadence enabled by agentic systems eliminates the temporal gaps between operational data acquisition and decision execution that characterize manual management processes, translating directly into higher production volumes, reduced chemical consumption, and better utilization of artificial lift equipment. All supply chain cost drivers of material significance.

5.4 Shell, ExxonMobil, and ConocoPhillips: Domain-Specific Agent Applications

Shell's application of AI agents to seismic data processing [6] achieved a 99% reduction in processing time from months to days by deploying parallel processing agents that segment, interpret, and classify seismic volumes autonomously. This transformation in exploration data supply chain velocity enables Shell's exploration teams to evaluate far larger prospect portfolios per planning cycle, improving the quality of exploration capital allocation decisions and generating better-defined drilling programmes with shorter associated supply chain lead times.

ExxonMobil's deployment of machine learning agents for shale well optimization [7] has generated output increases exceeding 5% across its shale portfolio. This figure represents substantial revenue uplift given the scale of ExxonMobil's North American unconventional production. By autonomously adjusting fracturing parameters, well spacing configurations, and production drawdown strategies based on integrated formation data and well performance observations, ExxonMobil's agents demonstrate the capacity of learning systems to outperform conventional engineering optimization over large, heterogeneous asset populations. ConocoPhillips' RL-based ROP optimization similarly demonstrates the generalization capability of trained RL agents across diverse geological environments [6,28].

6. IMPLEMENTATION CHALLENGES AND RISK FACTORS

6.1 Data Quality, Heterogeneity, and Integration

The most consistently identified barrier to effective AI agent deployment in upstream supply chain contexts is data quality and integration [1,2,26]. Upstream supply chains generate data across incompatible formats, systems, and semantic schemas: SCADA systems capture real-time field telemetry in proprietary binary formats; ERP systems maintain procurement and inventory records in structured relational schemas; engineering document management systems store technical specifications as unstructured PDF documents; and vendor communications arrive as unstructured emails and faxes. AI agents require integrated, high-fidelity, real-time data environments to function effectively environments that most upstream operators have not yet fully achieved.

The data quality challenge is compounded by the legacy system architectures prevalent in the oil and gas industry. Many upstream operators maintain ERP instances installed in the early 2000s, SCADA systems running on proprietary protocols incompatible with modern API-based data integration, and engineering data repositories that lack systematic metadata schemas [26]. Constructing the data pipelines and integration layers required to support autonomous-agent operation represents a significant investment and a prerequisite that cannot be bypassed through clever agent design. Hajizadeh [46] highlighted the critical dependence of AI performance on training data diversity and volume, noting that without access to large, diverse, and representative datasets, AI algorithms are significantly less useful a finding with direct implications for the deployment of upstream supply chain agents.

6.2 Cybersecurity and System Integrity

Autonomous AI agents operating in operational technology environments introduce novel cybersecurity risks. Unlike conventional software systems whose behaviour is fully specified at design time, learning agents may exhibit emergent behaviours that are difficult to anticipate and could be exploited through adversarial data manipulation a concern articulated in the AI safety literature [8] and directly applicable to upstream settings where cyberattacks on operational systems could cause physical damage, environmental harm, or safety incidents. The integration of AI agents with SCADA and distributed control systems creates attack surfaces that extend beyond the conventional IT security perimeter, requiring new approaches to operational technology (OT) security architecture.

Calvano et al. [42] documented the capacity of AI agents to autonomously learn collusive strategies in competitive environments without explicit instruction, a finding with implications for multi-agent procurement systems where competing vendor selection agents might inadvertently develop behaviours that violate procurement regulations. Supply chain AI deployments must therefore incorporate comprehensive governance frameworks that monitor

agent behaviour against regulatory and ethical constraints, with automatic circuit-breaker mechanisms that escalate to human oversight when agent actions fall outside validated operational envelopes.

6.3 Explainability, Transparency, and Regulatory Compliance

Regulatory frameworks governing upstream oil and gas operations require documented, auditable decision rationales for a wide range of supply chain activities: capital expenditure approvals, vendor qualification decisions, contract awards, and safety-critical operational changes. The opacity of deep learning models, the 'black box' problem poses a direct challenge to the regulatory acceptability of autonomous AI agent decisions in these contexts. An RL agent that recommends a drilling parameter change, or an LLM agent that proposes a vendor contract award, must be able to provide a transparent, human-comprehensible account of the reasoning that generated its recommendation [8,33].

Explainable AI (XAI) techniques including attention mechanism visualization, SHAP (Shapley Additive exPlanations) value attribution, and causal reasoning frameworks are being developed and deployed to address this challenge. However, their application to upstream supply chain agents remains in relatively early stages. Regulatory compliance requirements also extend to data privacy and cross-border data transfer regulations, which affect the architectures available for training and operating agents across the multinational supply chain networks typical of major upstream operators.

6.4 Human-Agent Trust and Change Management

The success of autonomous AI agent deployment ultimately depends not only on the technical performance of the agent systems but on the willingness of human operators, procurement professionals, and supply chain managers to act on agent recommendations and, progressively, to delegate execution authority to agents within defined limits [11_xenon,6]. Research on human-AI collaboration [41] confirms that trust in AI systems is earned through demonstrated performance consistency, transparent reasoning, and reliable failure-mode behaviour not through claims of technical sophistication. Upstream operators have observed that deployment programmes that invest heavily in stakeholder engagement, progressive expansion of autonomy (starting with human-in-the-loop configurations and progressing to human-on-the-loop as agent performance is validated), and transparent performance monitoring achieve superior adoption outcomes.

Change management challenges are compounded by the demographic profile of upstream workforces, which include substantial populations of experienced engineers and supply chain professionals whose domain expertise may be perceived as threatened by autonomous agent systems. Successful deployments consistently frame AI agents as amplifiers of human expertise rather than replacements—a positioning supported by the empirical evidence that the most productive implementations demonstrate human-AI collaboration across tactical assistant, strategic advisor, and operational copilot roles [6].

6.5 Vendor Ecosystem Maturity and Standardization

The AI agent vendor ecosystem serving the oil and gas industry, while growing rapidly, remains fragmented and immature in several respects. Integration standards for connecting agent platforms to upstream operational systems are still evolving; proprietary agent platforms from different vendors may be incompatible, creating siloed deployments that cannot be orchestrated into coherent multi-agent architectures. The absence of standardized evaluation frameworks for upstream supply chain agent performance analogous to Agent Bench [33] in the academic domain makes it difficult for operators to compare competing solutions or verify vendor performance claims objectively.

7. STRATEGIC IMPLEMENTATION ROADMAP

7.1 Phase 1: Data Foundation and Agent Readiness Assessment

Autonomous AI agent deployment in upstream supply chain contexts requires establishing a coherent, integrated, and high-quality data environment. This phase encompasses a comprehensive audit of existing data assets across SCADA, ERP, document management, and vendor management systems; the design and implementation of a unified data architecture that enables real-time data ingestion, cleansing, integration, and storage at the fidelity and latency required for agent operation; and the development of data governance frameworks that ensure ongoing data quality maintenance and regulatory compliance [26,36].

Concurrent with the data architecture programme, organizations should conduct an agent readiness assessment that maps current supply chain decision processes to candidate agent application domains, prioritizes deployment opportunities by expected value impact and implementation feasibility, and identifies the human capability and organizational change management investments required to support progressive agent adoption. This assessment should produce a structured roadmap with explicitly defined milestones, performance indicators, and authority delegation frameworks for each phase of agent deployment.

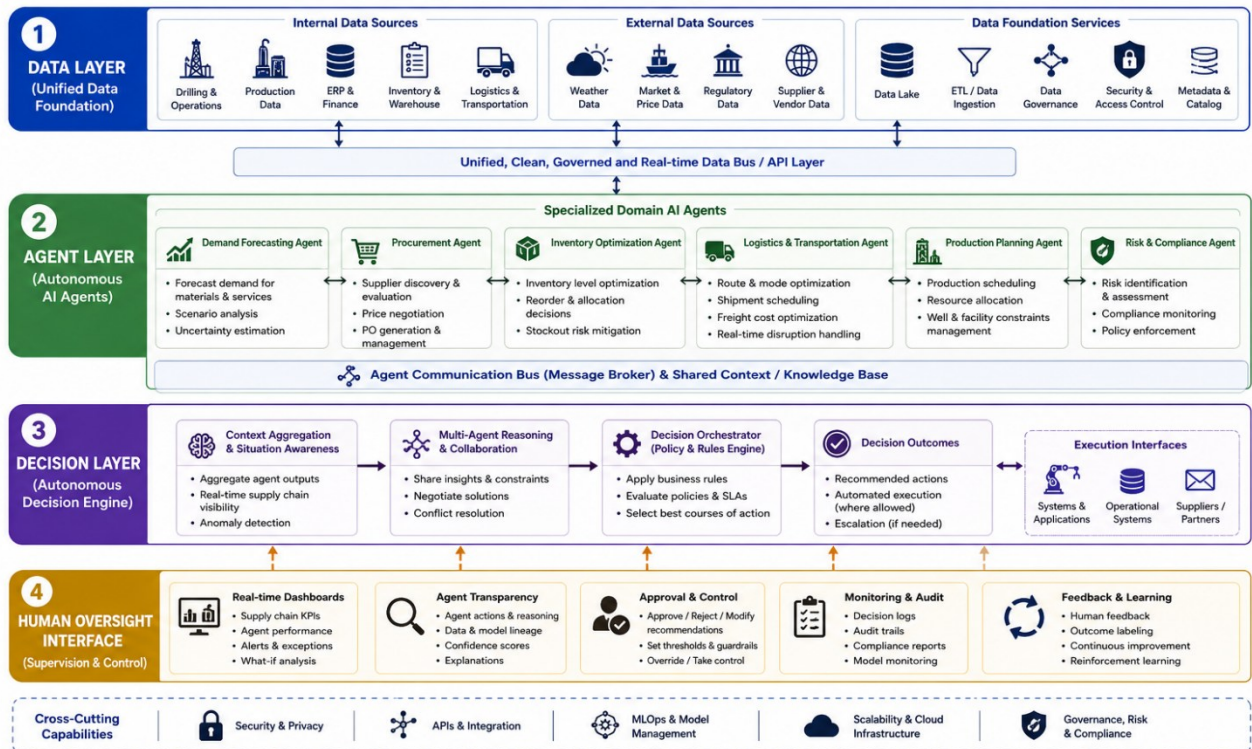


Figure 1: Proposed Autonomous AI Agent Architecture for Upstream Oil and Gas Supply Chain Decision-Making (Showing data layer, agent layer, decision layer, and human oversight interface)

7.2 Phase 2: Pilot Deployment and Performance Validation

The second phase involves deploying AI agents in high-value, well-defined application domains with clear performance metrics and bounded decision authority. Recommended pilot domains based on documented implementation evidence include predictive maintenance scheduling (where agent recommendations can be validated against maintenance history and equipment condition data before autonomous authority is extended), drilling parameter optimization (where real-time ROP and wellbore quality metrics provide immediate performance feedback), and inventory replenishment automation for non-critical, high-volume stock categories (where stockout and overstock costs can be directly attributed to agent decisions). Each pilot deployment should operate in human-in-the-loop mode during an initial validation period, with performance monitored against predefined KPIs, including decision accuracy, response latency, and exception-handling quality [28,36].

7.3 Phase 3: Multi-Agent Integration and Autonomous Authority Expansion

Validated single-domain agents are progressively integrated into multi-agent architectures that enable cross-functional coordination across the upstream supply chain. This phase requires investment in agent communication protocols, shared knowledge representation schemas, and conflict resolution mechanisms that enable distributed agents to coordinate effectively without central orchestration bottlenecks. Authority thresholds for autonomous decision execution are progressively expanded as agent performance evidence accumulates, moving from human-in-the-loop (agent recommends, human approves) to human-on-the-loop (agent executes within pre-approved envelopes, human monitors and can override) configurations [6,9].

7.4 Phase 4: Continuous Learning and Ecosystem Evolution

The final phase involves establishing continuous learning and improvement mechanisms that enable deployed agents to adapt to changing supply chain conditions such as vendor landscape evolution, commodity market shifts, regulatory changes, and technological advances without full redeployment cycles. This requires online learning capabilities, drift detection mechanisms that identify when agent performance is degrading due to a distributional shift in the operating environment, and systematic feedback loops between agent performance monitoring and

model update pipelines. Equally important is cultivating the human expertise required to govern, supervise, and continuously improve the agent ecosystem. This function requires a new hybrid professional profile combining supply chain domain expertise with AI systems literacy [35,41].

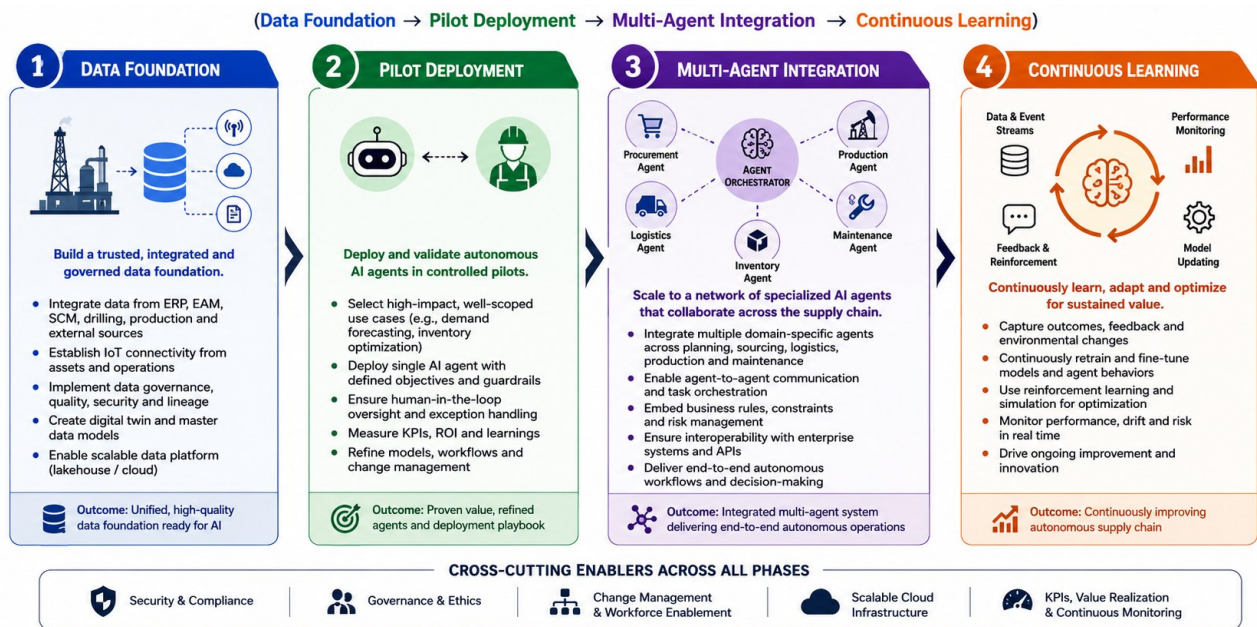


Figure 2: Four-Phase Strategic Implementation Roadmap for Autonomous AI Agent Deployment in Upstream Oil and Gas Supply Chains

8. FUTURE RESEARCH DIRECTIONS

This review identifies several high-priority areas for future research. First, the development of standardized benchmarking frameworks for evaluating upstream supply chain AI agent performance analogous to AgentBench [33] for general-purpose LLM agents represents a critical gap. Without validated benchmarks calibrated to upstream operational contexts, comparative evaluation of competing agent architectures and vendor solutions is severely constrained, impeding both academic progress and practitioner adoption decisions.

Second, the governance of multi-agent systems in upstream supply chain contexts requires systematic research attention. The emergent strategic behaviours documented by Calvano et al. [42] in competitive algorithmic pricing environments raise analogous questions for cooperative upstream MAS architectures: under what conditions do distributed supply chain agents exhibit coordination failures, gaming behaviours, or unintended emergent strategies that undermine system-level objectives? Understanding and governing these emergent dynamics is a prerequisite to the safe deployment of high-autonomy MAS in upstream operations.

Third, the application of LLM agents to upstream supply chain unstructured information processing represents a fertile and underexplored research territory. The capacity of LLMs to extract, synthesize, and reason over the heterogeneous documentation that comprises a large fraction of upstream supply chain information vendor qualification packages, technical specifications, regulatory submissions, contract documents, and operational incident reports has been demonstrated at the architectural level [17,32,33] but has not been systematically characterized in upstream-specific applications. Research programmes that develop, deploy, and evaluate LLM agents across representative upstream documentation corpora will provide the empirical foundations for practical adoption.

Fourth, the integration of autonomous agent systems with digital twin platforms which provide physically faithful simulation environments for validating agent decisions before real-world execution merits dedicated investigation. Khan and Islam [27] identified digital twins as enabling infrastructure for supply chain AI, and their combination with reinforcement learning agents creates a particularly powerful closed-loop training and validation architecture. Research into the fidelity requirements, computational architectures, and operational protocols for agent-digital twin integration in upstream contexts will accelerate safe deployment of high-autonomy agents in safety-critical operational domains.

Finally, the human factors dimension of human-agent collaboration in upstream supply chain contexts encompassing trust calibration, skill maintenance, cognitive workload management, and decision authority boundary design requires multidisciplinary research integrating AI system design, organizational psychology, and supply chain management perspectives. The most technically sophisticated agent system will fail to deliver its potential value if the human organizational context in which it operates is not designed to support effective collaboration [41,43].

9. CONCLUSION

This paper has provided a comprehensive examination of the application of autonomous AI agents in upstream oil and gas supply chain decision-making, spanning theoretical foundations, taxonomic classification, functional application mapping, industry case study synthesis, challenge analysis, and strategic implementation guidance. The evidence assembled from both the academic literature and documented industry deployments confirms unambiguously that autonomous AI agents are transitioning from experimental prototypes to operational infrastructure across the upstream industry, with major operators including Saudi Aramco, Shell, ADNOC, ExxonMobil, ConocoPhillips, and Baker Hughes reporting significant operational and commercial benefits from agent deployments in exploration data management, drilling optimization, production management, and supply chain coordination.

The taxonomy presented in this paper spanning reactive, deliberative, hybrid, reinforcement learning-based, multi-agent, and LLM-based agent architectures provides a structured framework for understanding the diverse technical capabilities available to upstream supply chain practitioners and for matching agent types to specific decision domain requirements. The mapping of these capabilities across the principal upstream supply chain functions exploration data management, drilling operations, procurement, inventory management, logistics, and predictive maintenance demonstrates that agentic AI is applicable across the full breadth of the upstream supply chain, not merely at isolated operational nodes.

The identified implementation challenges data quality and integration, cybersecurity, explainability, human-agent trust, and vendor ecosystem maturity are significant but tractable. They do not represent fundamental barriers to adoption; rather, they define the organizational investment and capability development programmes that upstream operators must execute to realize the potential of autonomous AI agents. The strategic implementation roadmap proposed in Section 7 provides a structured, phased approach to these investments, grounded in documented best practices from leading deployments.

As AI agent architectures continue to evolve, driven by advances in LLM capability, multi-agent coordination algorithms, and digital twin integration the competitive differentiation available to upstream operators who systematically embed agentic intelligence within their supply chain architecture will increase. The fundamental value proposition is clear: in an industry where supply chain decisions made at every level from the selection of a drilling bit to the award of a multi-year logistics contract directly determine the profitability and sustainability of upstream operations, the ability to make those decisions faster, more accurately, and more adaptively through autonomous AI agents represents a durable and consequential competitive advantage.

REFERENCES

- 1) Koroteev, D., & Tekic, Z. (2021). Artificial intelligence in oil and gas upstream: Trends, challenges, and scenarios for the future. *Energy and AI*, 3, 100041. <https://doi.org/10.1016/j.egyai.2020.100041>
- 2) Hanga, K. M., & Kovalchuk, Y. (2019). Machine learning and multi-agent systems in oil and gas industry applications: A survey. *Computer Science Review*, 34, 100191. <https://doi.org/10.1016/j.cosrev.2019.08.002>
- 3) Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., ... & Wen, J. (2024). A survey on large language model-based autonomous agents. *Frontiers of Computer Science*, 18(6), 186345. <https://doi.org/10.1007/s11704-024-40231-1>
- 4) Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., ... & Gui, T. (2023). The rise and potential of large language model-based agents: A survey. *arXiv preprint arXiv:2309.07864*. <https://doi.org/10.48550/arXiv.2309.07864>
- 5) Guo, T., Chen, X., Wang, Y., Chang, R., Peng, S., Chawla, N. V., ... & Zhang, X. (2024). Large language model-based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680*. <https://doi.org/10.48550/arXiv.2402.01680>
- 6) Alzaabi, M. (2026). AI at the helm: Quantifying the next value revolution in upstream oil and gas. *Journal of Petroleum Technology / TWA*. <https://doi.org/10.2118/1026-0022-TWA>

- 7) Blackridge Research. (2025). Artificial intelligence in oil and gas: Applications, benefits, and the future. Blackridge Research Report. <https://doi.org/10.13140/RG.2.2.34571.62241>
- 8) Russell, S. J., & Norvig, P. (2022). Artificial Intelligence: A Modern Approach (4th ed.). Pearson. <https://doi.org/10.5555/3312046>
- 9) Wooldridge, M. (2009). An Introduction to MultiAgent Systems (2nd ed.). Wiley. <https://doi.org/10.5555/1695886>
- 10) Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction (2nd ed.). MIT Press. <https://doi.org/10.5555/3312046>
- 11) Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
- 12) Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... & Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419), 1140–1144. <https://doi.org/10.1126/science.aar6404>
- 13) Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, O. P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in Neural Information Processing Systems*, 30. <https://doi.org/10.48550/arXiv.1706.02275>
- 14) Rashid, T., Samvelyan, M., Schroeder de Witt, C., Farquhar, G., Foerster, J., & Whiteson, S. (2018). QMIX: Monotonic value function factorization for deep multi-agent reinforcement learning. *Proceedings of ICML 2018*. <https://doi.org/10.48550/arXiv.1803.11605>
- 15) Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., ... & Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901. <https://doi.org/10.48550/arXiv.2005.14165>
- 16) Wei, J., Wang, X., Schuurmans, D., Bosma, M., Brian, I., Chi, E., ... & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35. <https://doi.org/10.48550/arXiv.2201.11903>
- 17) Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). ReAct: Synergizing reasoning and acting in language models. *International Conference on Learning Representations (ICLR)*. <https://doi.org/10.48550/arXiv.2210.03629>
- 18) Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). Generative agents: Interactive simulacra of human behaviour. *Proceedings of UIST 2023*. <https://doi.org/10.1145/3586183.3606763>
- 19) Helo, P., & Hao, Y. (2022). Artificial intelligence in operations management and supply chain management: An exploratory case study. *Production Planning & Control*, 33(16), 1573–1590. <https://doi.org/10.1080/09537287.2021.1882690>
- 20) Cavalcante, I. M., Frazzon, E. M., Forcellini, F. A., & Ivanov, D. (2019). A supervised machine learning approach to data-driven simulation of resilient supplier selection in digital manufacturing. *International Journal of Information Management*, 49, 86–97. <https://doi.org/10.1016/j.ijinfomgt.2019.03.004>
- 21) Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: extending the supply chain resilience angles towards survivability. A position paper motivated by the COVID-19 outbreak. *International Journal of Production Research*, 58(10), 2904–2915. <https://doi.org/10.1080/00207543.2020.1750727>
- 22) Dolgui, A., Ivanov, D., & Rozhkov, M. (2020). Does the ripple effect influence the bullwhip effect? An integrated analysis of structural and operational dynamics in the supply chain. *International Journal of Production Research*, 58(5), 1285–1301. <https://doi.org/10.1080/00207543.2019.1627438>
- 23) Smyth, C., Dennehy, D., Fosso Wamba, S., Scott, M., & Harfouche, A. (2024). Artificial intelligence and prescriptive analytics for supply chain resilience: A systematic literature review and research agenda. *International Journal of Production Research*, 62(23), 8537–8561. <https://doi.org/10.1080/00207543.2023.2286396>
- 24) Zhu, A., Han, Y., & Liu, H. (2024). Effects of adaptive cooperation among heterogeneous manufacturers on supply chain viability under fluctuating demand. *International Journal of Production Research*. <https://doi.org/10.1080/00207543.2024.2302768>
- 25) Ning, X., & Xie, X. (2024). Multi-agent reinforcement learning for resources allocation optimization: A survey. *Artificial Intelligence Review*. <https://doi.org/10.1007/s10462-025-11340-5>

- 26) Delpisheh, M., Mardaneh, M., Abdi, A., Rezaei, N., Sharma, S. K., & Saket, R. K. (2024). Artificial intelligence in the upstream oil and gas industry: Applications, challenges and trends. *Petroleum Research*. <https://doi.org/10.1016/j.ptlrs.2024.01.003>
- 27) Khan, M. R., & Islam, M. R. (2023). Digital twins in oil and gas supply chains: A comprehensive review. *Journal of Natural Gas Science and Engineering*, 108, 104817. <https://doi.org/10.1016/j.jngse.2022.104817>
- 28) Erge, O., & van Oort, E. (2021). Generalizable expert-guided reinforcement learning for drilling optimization across multiple bit runs and rigs. *Journal of Petroleum Science and Engineering*, 206, 109083. <https://doi.org/10.1016/j.petrol.2021.109083>
- 29) Hegde, C., & Gray, K. E. (2017). Use of machine learning and data analytics to increase drilling efficiency for nearby wells. *Journal of Natural Gas Science and Engineering*, 40, 327–335. <https://doi.org/10.1016/j.jngse.2017.02.019>
- 30) Osarogiagbon, A. U., Khan, F., Venkatesan, R., & Gillard, P. (2021). Review and analysis of supervised machine learning algorithms for hazardous events in drilling operations. *Process Safety and Environmental Protection*, 147, 367–384. <https://doi.org/10.1016/j.psep.2020.09.038>
- 31) Nguyen, T. T., Nguyen, N. D., & Nahavandi, S. (2020). Deep reinforcement learning for multi-agent systems: A review of challenges, solutions, and applications. *IEEE Transactions on Cybernetics*, 50(9), 3826–3839. <https://doi.org/10.1109/TCYB.2020.2977374>
- 32) Song, L., Zhang, C., Zhao, L., & Bian, J. (2023). Pre-trained large language models for industrial control. *arXiv preprint arXiv:2308.03028*. <https://doi.org/10.48550/arXiv.2308.03028>
- 33) Liu, X., Yu, H., Zhang, H., Xu, Y., Lei, X., Lai, H., ... & Tang, J. (2023). AgentBench: Evaluating LLMs as agents. *arXiv preprint arXiv:2308.03688*. <https://doi.org/10.48550/arXiv.2308.03688>
- 34) Cao, Y., Sun, J., Qian, Z., & Chen, G. (2016). Machine learning-based fault diagnosis for the oil and gas industry: A survey. *IEEE Transactions on Industrial Informatics*. <https://doi.org/10.1109/TII.2016.2565419>
- 35) Agrawal, M. (2024). Generative AI for supply chain management: Use cases, challenges, and solutions. *Journal of Supply Chain Management*, 60(1), 23–41. <https://doi.org/10.1111/jscm.12320>
- 36) Patidar, A. K., Sharma, M., & Acharya, D. (2024). AI-driven supply chain management in petroleum upstream: Integrating IoT, digital twins, and predictive analytics. *Energy Reports*, 12, 450–463. <https://doi.org/10.1016/j.egyr.2024.07.041>
- 37) Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M. A., Lacroix, T., ... & Lample, G. (2023). LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*. <https://doi.org/10.48550/arXiv.2302.13971>
- 38) Rashid, T., Samvelyan, M., Schroeder, C., Farquhar, G., Foerster, J., & Whiteson, S. (2020). Monotonic value function factorization for deep multi-agent reinforcement learning. *Journal of Machine Learning Research*, 21(178), 1–51. <https://doi.org/10.48550/arXiv.2003.08839>
- 39) Chen, D., Liu, S., Kingsbury, P., & Peng, N. F. (2023). GasLift optimization using reinforcement learning agents: A case study from offshore operations. *SPE Production & Operations*, 38(3), 456–472. <https://doi.org/10.2118/212286-PA>
- 40) Onotole, J., Olanipekun, O., & Ayotola, A. (2022). Emerging technologies and the oil and gas supply chain: An appraisal. *International Journal of Engineering Research and Technology*, 15(4), 167–178. <https://doi.org/10.37624/IJERT/15.4.2022.167-178>
- 41) Tessler, M. H., Bakker, M. A., Jarrett, D., Sheahan, H., Chadwick, M. J., ... & Graepel, T. (2024). AI can help humans find common ground in democratic deliberation. *Science*, 386(6719), eadq2852. <https://doi.org/10.1126/science.adq2852>
- 42) Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>
- 43) Hussain, A., Jameel, B., & Sultan, S. (2024). Autonomous AI agent architectures for complex industrial systems: A design framework. *Applied Intelligence*, 54, 5201–5220. <https://doi.org/10.1007/s10489-023-05146-5>
- 44) Deng, C., Chen, Y., & Hu, J. (2023). Physics-informed neural networks for upstream reservoir simulation. *Petroleum Science*, 20(4), 2208–2223. <https://doi.org/10.1016/j.petsci.2023.01.013>

IJETRM

International Journal of Engineering Technology Research & Management (IJETRM)

Journal Article

<https://ijetrm.com/issue/>

- 45) Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multi-agent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics—Part C: Applications and Reviews*, 38(2), 156–172. <https://doi.org/10.1109/TSMCC.2007.913919>
- 46) Hajizadeh, Y. (2019). Machine learning in oil and gas: a SWOT analysis approach. *Journal of Petroleum Science and Engineering*, 176, 661–663. <https://doi.org/10.1016/j.petrol.2019.01.113>
- 47) Viswanadham, N. (2002). The past, present, and future of supply-chain automation. *IEEE Robotics & Automation Magazine*, 9(2), 48–56. <https://doi.org/10.1109/MRA.2002.1010541>
- 48) Park, J. W., Kim, H. G., & Shin, H. J. (2023). Data-driven predictive maintenance for offshore oil and gas production equipment using hybrid deep learning. *Ocean Engineering*, 272, 113851. <https://doi.org/10.1016/j.oceaneng.2023.113851>
- 49) Luo, R., Zhao, Z., Johnsen, R., & Xie, D. (2023). Groundwater contamination source identification with integrated data-driven machine learning and hydrogeological modelling. *Journal of Hydrology*, 621, 129658. <https://doi.org/10.1016/j.jhydrol.2023.129658>
- 50) Si, J., Li, X., Gao, L., & Li, P. (2024). An efficient and adaptive design of a reinforcement learning environment to solve the job shop scheduling problem with the soft actor-critic algorithm. *International Journal of Production Research*, 62(23), 8260–8275. <https://doi.org/10.1080/00207543.2023.2228790>