Governance, Accountability, and Regulatory Compliance of AI-Driven Autonomous Drilling Control Systems in Critical Energy Infrastructure

Volodymyr Hnatushenko[1,†], Volodymyr Khomenko[1,*,†], Oleksandr Pashchenko[1,†] and Serhii Hryshchak[1,†]

[1] *Dnipro University of Technology, Dmytra Yavornytskoho Ave 19, 49005, Dnipro, Ukraine*

Abstract

The rapid adoption of AI-driven autonomous drilling control systems (ADCS) in oil, gas, and geothermal operations positions them as critical energy infrastructure. However, real-time autonomous decision-making creates legal and ethical gaps that existing frameworks do not address. For example, when an AI agent causes a subsurface collision or blowout, it remains unclear whether the operator, the AI vendor, or the data provider bears liability. This paper analyzes three pillars: liability allocation, algorithmic transparency, and regulatory compliance. We propose a multi-layered responsibility matrix that distinguishes operator duties (safe envelope and fallback), vendor liability (training data and validation), and integrator responsibility (sensor integrity and handshake protocols). For transparency, we argue for mandated recordable AI decision logs – a "drillable black box" – that capture inputs, actions, confidence scores, and overrides, while reconciling trade secrecy with critical infrastructure audit requirements. Compliance analysis reveals conflicts between real-time control and logging overhead, as well as mismatches with offshore safety directives and cybersecurity standards (e.g., IEC 62443). Without regulatory updates, autonomous drilling creates uninsured risks, including unallocated liability for cascading failures. We conclude by recommending hybrid human-AI supervision models (asymmetric authority with right-to-intervene) and certification standards for drilling AI, inspired by aviation practices, to ensure safe, transparent, and accountable deployment in critical infrastructure.

Keywords

critical infrastructure; AI liability; algorithmic transparency; regulatory compliance; well control; ethical AI

1. Introduction

Autonomous drilling control systems (ADCS) represent a paradigm shift in subsurface extraction operations, evolving from simple automated pipe rotation and anti-stick-slip controllers to fully closed-loop geosteering and real-time parameter optimization [1, 2]. Modern ADCS integrate downhole telemetry, surface sensors, and machine learning models that continuously adjust weight on bit, rotational speed, flow rate, and trajectory without human intervention. These systems are deployed on drilling rigs – both offshore platforms and on land – that are unequivocally part of critical energy infrastructure, where uninterrupted and safe operation is essential for economic stability, energy security, and environmental protection [3, 4]. The promise of ADCS is substantial: field studies indicate increased rate of penetration by 15-30%, reduction of non-productive time due to fewer drilling dysfunctions, and, most importantly, lower human exposure to hazardous environments such as high-pressure zones, flammable gases, and heavy machinery [5, 6]. By shifting routine and complex control decisions from human drillers to algorithms, proponents argue that ADCS can enhance consistency, reduce fatigue-related errors, and operate at speeds impossible for

[1] *RS-2026: Resilient Systems: Secure Digital Technologies and Critical Infrastructure, June 30, 2026, Drohobych, Ukraine*

* Corresponding author.

† These authors contributed equally.

hnatushenko.v.v@nmu.one (V. Hnatushenko); homenko.v.l@nmu.one (V. Khomenko); pashchenko.o.a@nmu.one (O. Pashchenko); hryshchak.s.v@nmu.one (S. Hryshchak)

0000-0003-3140-3788 (V. Hnatushenko); 0000-0002-3607-5106 (V. Khomenko); 0000-0003-3296-996X (O. Pashchenko); 0000-0002-6302-1322 (S. Hryshchak)

manual control. Similar trends toward autonomous decision support and intelligent control have already been demonstrated in critical industrial and energy systems, including intelligent electricity load forecasting and smart regulation of mine drainage facilities, where machine learning models are increasingly entrusted with operational decisions under safety constraints [4, 7].

However, the same characteristics that make ADCS attractive also create profound legal and ethical challenges [8, 9]. Because these systems operate on critical infrastructure, any failure – a subsurface collision with another well, a blowout caused by misinterpreted pore pressure, or a stuck pipe event due to an erroneous optimization policy – can escalate into environmental disasters (e.g., uncontrolled hydrocarbon releases), catastrophic asset loss (a rig worth hundreds of millions of dollars), or fatalities among crew and nearby populations. Unlike autonomous vehicles, which operate in open but observable environments, drilling autonomy confronts deep subsurface uncertainty, delayed feedback loops (it may take hours to know that a previous decision led to a wellbore collapse), and irreversible actions with long-consequence chains. The core thesis of this paper is that current legal and ethical frameworks are fundamentally unprepared for autonomous decision-making in subsurface real-time control [10]. Existing liability doctrines assume a human operator who directly controls or supervises every actionable step; transparency requirements rely on post-incident reconstruction of deterministic actions; and critical infrastructure regulations mandate "competent persons" in charge without defining what that means when an AI agent executes sub-second adjustments [11, 12]. This introduction sets the stage for a systematic analysis of three interconnected gaps: liability allocation among operators, AI vendors, and data providers; the tension between algorithmic transparency and trade secrecy under regulatory audit; and compliance mismatches with sector-specific rules such as offshore safety directives and industrial cybersecurity standards [13, 14]. By grounding the discussion in concrete drilling scenarios, the paper aims to provide actionable recommendations for regulators, industry, and legal practitioners before an autonomous-driven disaster forces reactive and possibly ill-fitting legal changes.

2. Relevance of Research

This research directly aligns with several core themes of the conference. Most centrally, it addresses ethical considerations and regulatory frameworks for the use of artificial intelligence in critical infrastructure, examining how autonomous drilling control systems challenge existing notions of accountability, fairness, and safety [15, 16]. Second, it contributes to the conference topic on integration of AI into real-time systems for managing critical facilities, using drilling as a concrete, high-stakes use case where control loops operate at sub-second intervals and decisions have irreversible subsurface consequences [17]. Third, it speaks to resilience and security, because autonomous control fundamentally reshapes operational resilience – an AI agent that optimizes for short-term penetration rate may inadvertently reduce long-term wellbore stability or create new failure modes that human operators are not trained to recognize [18].

Recent research on intelligent decision-support systems in energy and industrial environments demonstrates that increasing levels of autonomy require corresponding mechanisms for accountability, auditability, and risk management. This challenge is especially relevant for critical infrastructure, where AI-driven forecasting and optimization systems directly influence operational safety and reliability [19].

Beyond conference alignment, the research is motivated by growing regulatory interest worldwide. The European Union's AI Act classifies certain AI systems used in critical infrastructure as high-risk, imposing requirements for risk management, data governance, transparency, and human oversight – yet drilling autonomy remains largely absent from sectoral guidance. Industry bodies such as the International Association of Drilling Contractors (IADC) and the International Association of Oil & Gas Producers (IOGP) have begun issuing guidelines on automated drilling, but these documents focus on technical reliability rather than legal liability or regulatory compliance [20, 21]. The practical relevance of this research is immediate and multi-stakeholder. Drilling contractors need clarity on how to allocate risk in contracts that include ADCS; oil and gas operators

must understand their duty of supervision under health, safety, and environment (HSE) regulations; AI vendors face potential product liability claims if their algorithms cause incidents; and insurance underwriters are currently unable to price premiums for autonomous drilling due to the absence of legal precedents and regulatory standards. Finally, this research fills a distinct academic gap [22]. The vast majority of literature on legal and ethical aspects of autonomous systems concerns autonomous vehicles (where the environment is observable and decisions are short-lived) or aviation (where certification pathways exist and human pilots remain in the loop) [23, 24]. Drilling autonomy is fundamentally different: subsurface uncertainty means the AI cannot directly "see" the state of the rock; decision consequences unfold over hours or days; and the cost of failure is often environmental catastrophe rather than simply property damage. By focusing on these unique characteristics, the paper provides a novel contribution that is both theoretically grounded and operationally urgent.

3. Statement of the Problem

The deployment of autonomous drilling control systems (ADCS) on critical infrastructure creates a set of interconnected legal, transparency, and compliance problems that existing frameworks fail to address. These problems can be grouped into three clusters, but they operate simultaneously in practice, amplifying each other.

Core legal problems. The first cluster concerns liability. When an ADCS causes a well control incident – for example, a blowout resulting from the AI misinterpreting formation pressure data – the current legal regime offers no clear answer as to whether the operator bears responsibility for inadequate supervision, the AI vendor for flawed algorithmic design, or the data provider for inaccurate sensor fusion or training data. This liability vacuum is exacerbated by a fundamental doctrinal ambiguity: should ADCS be treated as a "product" (triggering strict liability for defects) or as a continuously learning "service" (subject only to negligence standards)? The distinction has profound consequences for who can be sued and under what theory. Compounding the problem, existing drilling contracts – such as the standard model forms published by the International Association of Drilling Contractors (IADC) – assume a human driller in physical control. They contain no clauses addressing autonomous mode, allocation of algorithmic risk, or indemnification for AI-caused events. The following diagram (Fig. 1) illustrates the current fragmented responsibility landscape.

Transparency and explainability problems. The second cluster arises from the black-box nature of many ADCS implementations. Deep reinforcement learning policies that optimize drilling parameters – weight on bit (WOB), rotations per minute (RPM), flow rate, and steering commands – produce effective actions but lack causal explanations. When a regulator or court asks why the AI decided to drill through a predicted loss zone instead of setting a casing shoe, the system cannot provide a human-understandable justification. This failure has direct legal consequences: under critical infrastructure regulations, operators must be able to demonstrate that decisions were reasonable and based on available data. Without explainability, post-incident reconstruction becomes impossible, and regulators cannot audit whether the AI complied with accepted practices. Furthermore, vendors aggressively protect their models as trade secrets, refusing to disclose training data, network architectures, or weight parameters. While trade secrecy is legitimate for commercial software, it collides directly with the auditability requirements that apply to critical infrastructure. A drilling AI that cannot be inspected by an independent authority is, from a regulatory perspective, indistinguishable from an uncontrolled hazard.

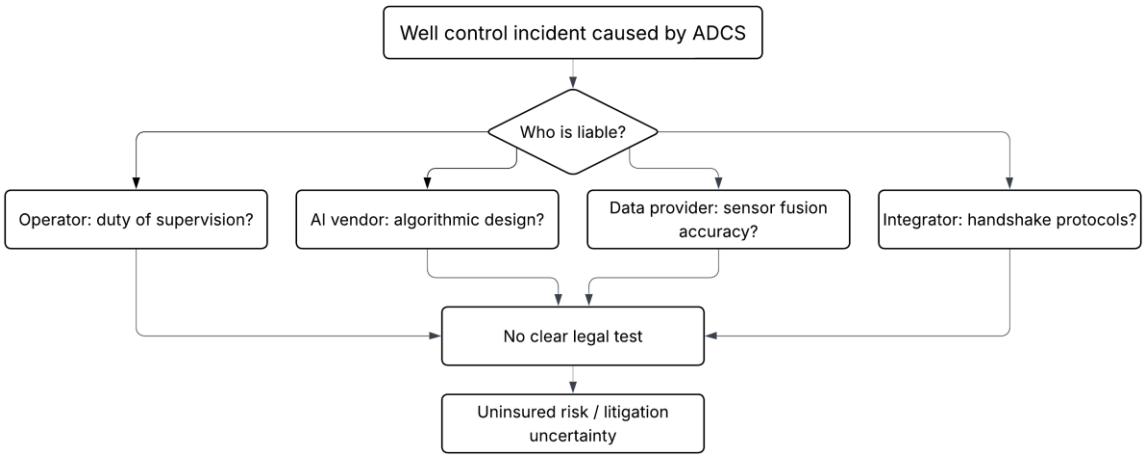

Figure 1: Current fragmented liability landscape for ADCS-caused well control incidents

Compliance conflicts. The third cluster involves specific conflicts between ADCS characteristics and existing regulations. First, cybersecurity: autonomous systems dramatically increase the attack surface. Similar concerns have been identified in critical communication infrastructures, where the integrity of autonomous decision-making depends on secure authentication, authorization, and protection against data manipulation. Emerging post-quantum authentication mechanisms and secure communication architectures provide useful models for future autonomous drilling environments [24,25]. An adversary could spoof downhole sensor readings (e.g., pressure or gamma ray logs) or perform reward hacking – manipulating the reinforcement learning environment to induce dangerous actions. Compliance with frameworks such as the EU NIS2 Directive and the IEC 62443 series for industrial control systems becomes ambiguous because those standards were designed for deterministic, human-supervised systems, not for learning agents that change behavior over time. Second, real-time constraints: effective drilling control requires sub-second loop closure. Mandating comprehensive transparency logs – recording every input, internal state, and output – would create storage and processing overhead that conflicts with low-latency requirements. This is not merely a technical trade-off but a legal one: if a regulator demands full logging and the operator omits it for performance reasons, the operator becomes non-compliant; if the operator implements it and the system fails to react in time, the operator may be liable for the resulting incident. The following formula expresses the compliance tension qualitatively:

$$\text{Compliance Risk} = f\left(\frac{\text{Logging Overhead}}{\text{Available Real-Time Budget}} \middle| \text{Regulatory Expectation of Completeness}\right) \quad (1)$$

When logging overhead approaches the real-time budget, the operator is forced to choose between two unacceptable outcomes. Third, human oversight requirements: nearly every critical infrastructure regulation (e.g., offshore safety cases, onshore HSE rules) demands that a "competent person" be in control of operations. But what does "control" mean when an AI agent adjusts drilling parameters every 200 milliseconds? Is a human who monitors a dashboard and can press an emergency stop button still "in control"? If the human overrides too late, does the legal responsibility shift? Regulators have not yet defined the necessary degree of human authority, latency, or intervention authority for autonomous drilling. The result is a legal gray zone where operators deploy ADCS without knowing whether they are compliant – a situation that is unsustainable for critical infrastructure protection.

4. Research Results

The analysis of legal and ethical gaps leads to four interconnected results: a liability framework that allocates responsibility across three actors, a transparency mechanism in the form of a mandated "drillable black box," a compliance roadmap for aligning autonomous drilling with existing critical infrastructure regulations, and an ethical matrix to guide deployment decisions. Each result is presented below as part of a coherent proposal.

Liability framework proposal. The current liability vacuum can be resolved through a trichotomy of responsibility that distinguishes the operator, the AI vendor, and the system integrator. The operator retains the duty to maintain a safe operational envelope – meaning that even in autonomous mode, the operator must define geomechanical and hydraulic boundaries (e.g., maximum equivalent circulating density, minimum allowable formation fracture pressure) that the ADCS cannot violate. The operator is also responsible for a fail-safe supervision architecture, including a human-accessible emergency stop that overrides any AI command, and a documented fallback plan for scenarios where the AI loses confidence or encounters conditions outside its training distribution. The AI vendor, in turn, bears liability for the representativeness of training data (e.g., whether the dataset includes rare but physically possible subsurface conditions), the coverage of simulator validation (i.e., testing across the entire operational design domain), and algorithmic stability (e.g., the absence of chaotic or bifurcating behaviors in response to small sensor perturbations). The integrator – the entity that embeds the AI into the rig's control system and data infrastructure – is liable for safe human-AI handshake protocols (ensuring that mode switches are unambiguous and logged) and for sensor integrity (detecting spoofing, drift, or failures before they corrupt the AI's inputs). A key legal innovation is the rebuttable presumption of operator control: if the ADCS takes an action that falls within its predefined Operational Design Domain (ODD) – a formally documented set of geological, operational, and environmental conditions for which the AI was certified – then liability is presumed to lie with the vendor or integrator. Conversely, if the AI acts outside its ODD, a presumption shifts to the operator, who should have either intervened or not activated autonomous mode. The following diagram (Fig. 2)summarizes the liability allocation.

For insurance, this framework implies that traditional well control policies (underwritten for human-controlled drilling) are insufficient. We recommend mandatory algorithmic error and omission (E&O) coverage for drilling AI vendors, with premiums tied to the diversity of training data and the frequency of offline policy audits.

Transparency mechanism – "drillable black box." To address the explainability deficit, we propose a mandated real-time logging system analogous to flight data recorders but adapted for AI decision-making. For every control cycle (e.g., every 200 milliseconds), the ADCS must record: (1) all AI input features derived from sensor fusion (e.g., downhole pressure, torque, gamma ray, vibration spectra); (2) the action actually selected (e.g., WOB = 15,000 lbs, RPM = 120) along with up to three alternative candidate actions that were considered but rejected; (3) the confidence score assigned by the AI to each candidate and a computed risk estimate (e.g., probability of exceeding a safe operating limit); and (4) any human override events, including the time delay between the AI's action and the override. The legal status of this log is critical: it must be admissible as evidence in civil and administrative proceedings, and it must be subject to regulatory inspection without any trade secrecy exemption. Vendors cannot refuse to disclose the log on the grounds that it contains proprietary algorithms – because the log records only inputs, outputs, and internal scores, not the full weight matrix. However, where even that level of detail is contested, critical infrastructure regulations should override trade secrecy claims in the interest of public safety. Technically, storing every control cycle for a multi-week drilling operation would generate petabytes of data. Therefore, we specify selective storage: the system continuously buffers the last 60 seconds of high-frequency data; upon detection of an anomalous event (e.g., pressure spike, unexpected formation change, human override, or AI confidence dropping below a threshold), the buffer is permanently saved. This balances forensic completeness with feasibility. The trade-off can be expressed as:

$$\text{Log Retention} = \begin{cases} \text{Continuous buffer: 60 sec,} & \text{normal operation} \\ \text{Permanent storage: triggered by anomaly,} & \text{event condition} \end{cases}. \quad (2)$$

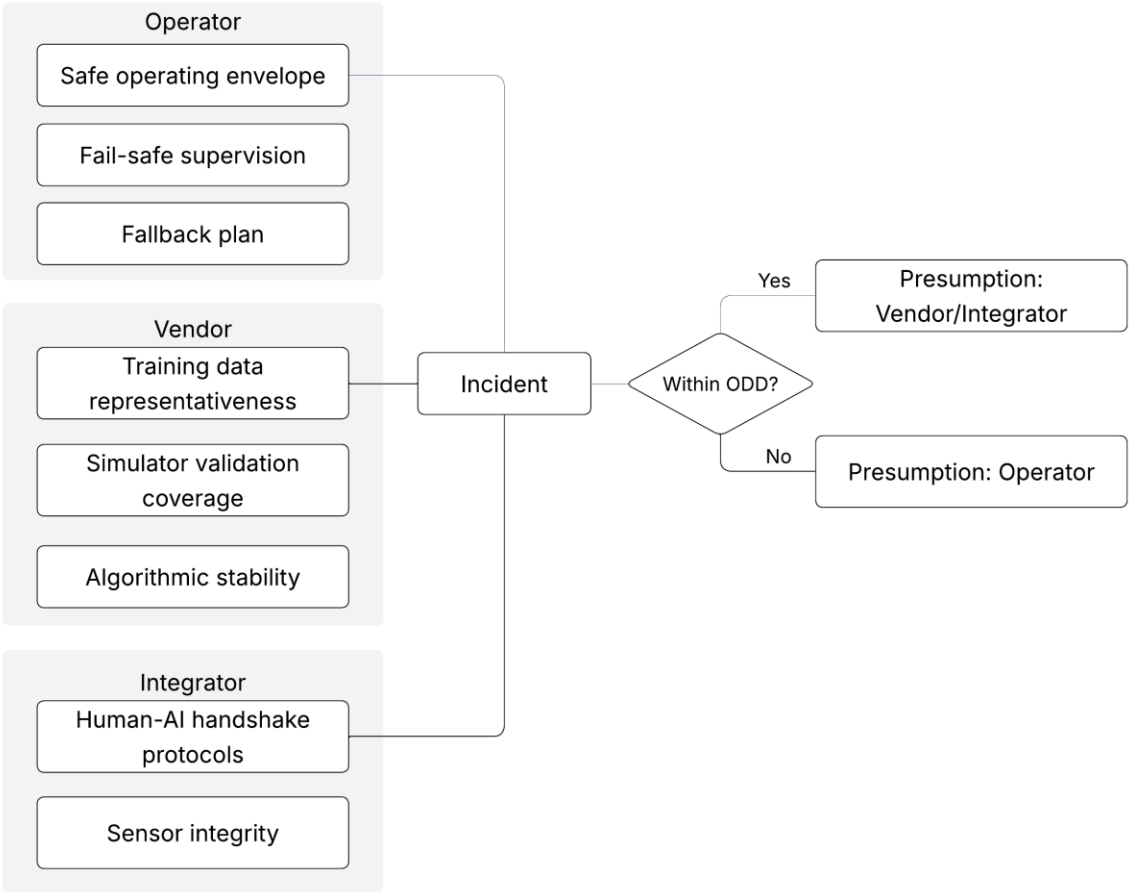

Figure 2: Trichotomy of responsibility and the rebuttable presumption of operator control based on Operational Design Domain (ODD)

Compliance roadmap for ADCS under existing critical infrastructure regulations. Mapping ADCS to the EU AI Act's high-risk categories (Annex III, point 2 on critical infrastructure operation) confirms that autonomous drilling systems fall squarely within the scope, requiring conformity assessments, risk management, and human oversight. However, sector-specific regulations such as API RP 98 (well control) currently contain no requirement for AI transparency or algorithmic auditing. The gap analysis reveals specific omissions: no definition of ODD for AI, no mandatory logging of AI decisions, and no requirement for periodic third-party evaluation of the autonomous control policy. To remedy this, we suggest three regulatory amendments. First, mandate that any ADCS deployed on critical infrastructure formally document its Operational Design Domain – including lithology types, pore pressure ranges, temperature limits, and allowable deviation from planned trajectory. Second, require periodic offline reverse engineering of black-box policies for regulatory audit. This means that at least once per six months, an independent auditor must extract the AI's decision function (e.g., via surrogate model approximation or input-output mapping) and verify that it does not produce unsafe actions within the ODD. Third, establish a certification body for drilling AI inspired by aviation's DO-178C (software considerations for airborne systems) but adapted to real-time subsurface control. Certification would include requirements for training data provenance, simulator fidelity, and closed-loop testing across edge cases. The following timeline (Fig. 2) illustrates a proposed compliance pathway for an operator deploying a new ADCS.

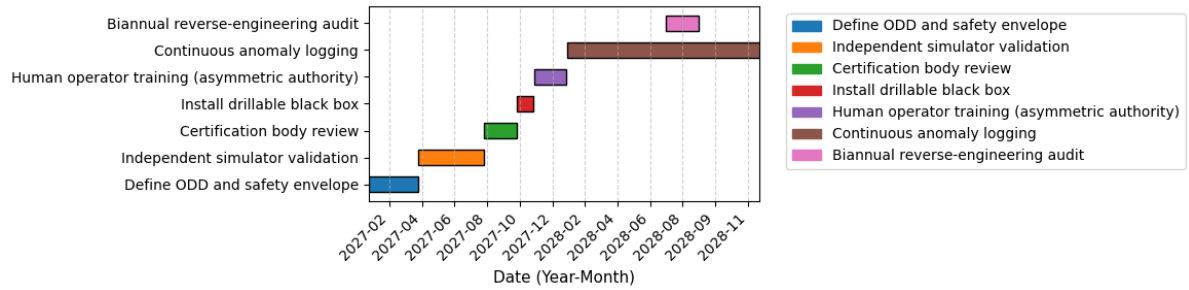

Figure 3: Proposed compliance pathway timeline for deploying a new autonomous drilling control system

Ethical matrix for autonomous drilling. Beyond legal compliance, four ethical principles must guide ADCS deployment. Beneficence requires that the system reduce health, safety, and environmental (HSE) exposure – but this benefit cannot override local community rights (e.g., prior informed consent for drilling near water sources) or environmental rights (e.g., protection of sensitive ecosystems). Non-maleficence imposes a higher standard: the AI must be demonstrably safer than a human driller operating under the same conditions. The burden of proof lies on the operator, who must provide statistical evidence (e.g., from simulator-based comparative trials) that the autonomous system reduces the probability of a blowout, loss of well control, or subsurface collision. Autonomy versus control demands that no full autonomy be permitted without continuous human right-to-intervene – an asymmetric authority where the human can always override the AI, but the AI cannot override the human. This implies that the ADCS must include a physical or logical "big red button" that disengages autonomous control and returns to a safe state (e.g., maintain current parameters or execute a pre-programmed shutdown sequence) within a latency not exceeding 500 milliseconds. Finally, justice prohibits shifting risk from operators to nearby populations or groundwater resources. For example, an AI policy that increases drilling rate to save operator time but raises the probability of a minor hydrocarbon seep into an aquifer would be unjust, because the benefit accrues to the operator while the risk falls on uninformed third parties. These ethical constraints are not merely aspirational: they can be encoded as constraints in the AI's reward function (e.g., penalizing any action that increases predicted environmental impact beyond a baseline). Together, the liability framework, transparency mechanism, compliance roadmap, and ethical matrix provide a coherent foundation for deploying autonomous drilling control systems that are legally defensible, regulatorily compliant, and ethically sound.

5. Conclusions

Autonomous drilling control systems operating on critical energy infrastructure remain legally and ethically underdefined, with liability gaps, transparency deficits, and regulatory mismatches that together create unacceptable risk of uninsured environmental disasters, asset loss, and fatalities. This paper has made four key contributions: first, a three-party liability allocation model distinguishing operator duties (safe envelope, fail-safe supervision, fallback plan), vendor liability (training data representativeness, simulator validation, algorithmic stability), and integrator responsibility (human-AI handshake protocols, sensor integrity), including a rebuttable presumption of operator control when the AI acts outside its Operational Design Domain. Second, a proposal for a mandated "drillable black box" logging system that records inputs, selected actions with alternatives, confidence scores, risk estimates, and human overrides, with selective storage of anomalous episodes and no trade secrecy exemption for regulatory inspection. Third, identification of specific conflicts between trade secrecy and audit requirements, as well as between real-time control latency and comprehensive logging overhead. Fourth, suggested regulatory updates for API RP 98 (mandating ODD, periodic offline reverse-engineering audits) and for AI Acts (clarifying high-risk status of drilling AI), alongside a certification body inspired by aviation DO-178C. The practical takeaway for

conference attendees is immediate: operators should not deploy ADCS without contractual AI vendor liability clauses and independent algorithmic auditing; regulators must act before a major autonomous-caused blowout forces reactive, ill-fitting rules. Future work will focus on empirical testing of the proposed transparency overhead on real-time drilling performance and the development of formal certification standards for drilling AI, including benchmark test suites for subsurface edge cases.

Acknowledgements

The authors thank the conference organizers and anonymous reviewers for their constructive feedback. We are also grateful to drilling engineers and legal experts from Dnipro University of Technology for their insights on real-world ADCS deployment challenges, though all conclusions remain solely those of the authors.

Declaration on Generative AI

During the preparation of this work, the author(s) used ChatGPT-4 (OpenAI) and Grammarly in order to: grammar and spelling check, style improvement, and rephrasing of selected sentences for clarity. After using these tools/services, the authors reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] M. C. Abadía, M. Goisauf, I. Hesso, R. Kayyali, M. Kogut-Czarkowska, R. Martínez, P. N. Jiménez, Societal, legal, and ethical aspects of trustworthy AI, in: Trustworthy AI in Cancer Imaging Research, Springer, Cham, 2025, pp. 3–21. doi:10.1007/978-3-031-89963-8_1.

[2] C. Cath, Governing artificial intelligence: Ethical, legal and technical opportunities and challenges, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 376 (2018) 20180080. doi:10.1098/rsta.2018.0080.

[3] M. Schedl, V. W. Anelli, E. Lex, Trustworthy recommender systems: Technical, ethical, legal, and regulatory perspectives, in: Proceedings of the 17th ACM Conference on Recommender Systems, ACM, New York, NY, 2023, pp. 1288–1290. doi:10.1145/3604915.3609497.

[4] F. N. Çiçin, S. Sayin, Understanding artificial intelligence along with legal and ethical issues, Journal of Oncological Sciences 10 (2024) 105–114. doi:10.37047/jos.2023-99598.

[5] J. Mbounkap, H. Djigo, T. Masheh, A. Nzooh, A. Kar, Technical and ethical perspectives on building intelligent legal assistants with LLM and RAG pipelines, in: Advances in Computational Intelligence and Robotics, IGI Global Scientific Publishing, Hershey, PA, 2026, pp. 411–468. doi:10.4018/979-8-3373-7685-1.ch013.

[6] Ratov, B., Pavlychenko, A., Kirin, R., Pashchenko, O., Khomenko, V., Tileuberdi, N., Kamyshatskyi, O., Sieriebriak, S., Seidaliyev, A., & Muratova, S. (2025). Using machine learning to model mechanical processes in mining: Theory, practice, and legal considerations. Engineered Science, 33, 1419. https://doi.org/10.30919/es1419.

[7] R. Kirin, B. Ratov, V. Khomenko, O. Pashchenko, S. Hryshchak, A. Shukmanova, L. Nurshakhanova, Artificial intelligence technologies in the mining industry: Economic and legal assessment, Mining of Mineral Deposits 20 (1) (2026) 125–141. doi:10.33271/mining20.01.123.

[8] R. Lath, R. Patwari, A. Aylani, D. Hajoary, Introduction to generative AI, in: Generative AI, Wiley, Hoboken, NJ, 2025, pp. 1–27. doi:10.1002/9781394302932.ch1.

[9] R. Dreyling, E. Jackson, T. Tammet, A. Labanava, I. Pappel, Social, legal, and technical considerations for machine learning and artificial intelligence systems in government, in: Proceedings of the 23rd International Conference on Enterprise Information Systems, SCITEPRESS, Setúbal, 2021, pp. 701–708. doi:10.5220/0010452907010708.

[10] G. Singh, V. Srivastava, S. Kumar, V. Bhatnagar, S. A. Dhondiyal, Understanding the ethical implications of generative AI: A multi-disciplinary perspective, in: 2024 5th International Conference on Electronics and Sustainable Communication Systems (ICESC), IEEE, Piscataway, NJ, 2024, pp. 1714–1719. doi:10.1109/icesc60852.2024.10689798.

[11] T. Haidemariam, A.-B. Gran, On the problems of training generative AI: Towards a hybrid approach combining technical and non-technical alignment strategies, AI & SOCIETY 41 (2025) 629–654. doi:10.1007/s00146-025-02445-0.

[12] M. Schedl, V. W. Anelli, E. Lex, Trustworthy user modeling and recommendation from technical and regulatory perspectives, in: Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization, ACM, New York, NY, 2024, pp. 17–19. doi:10.1145/3631700.3658522.

[13] V. Jain, A. Mitra, Ensuring compliance and ethical standards with generative AI in fintech, in: Generative Artificial Intelligence in Finance, Wiley, Hoboken, NJ, 2025, pp. 253–264. doi:10.1002/9781394271078.ch13.

[14] L. Marassi, A. E. Pascarella, G. Giacco, M. Zuccarini, S. Marrone, C. Sansone, D. Amitrano, M. Rigiroli, Artificial intelligence and voluntary carbon marketplaces: An analysis of the ethical and legal aspects, in: Proceedings of the 2023 Conference on Human Centered Artificial Intelligence: Education and Practice, ACM, New York, NY, 2023, pp. 53–53. doi:10.1145/3633083.3633225.

[15] O. Hniche, R. Saadane, Challenges and opportunities of integrating AI with IFRS in accounting systems insights from Morocco, in: Proceedings of the 6th International Conference on Networking, Intelligent Systems & Security, ACM, New York, NY, 2023, pp. 1–9. doi:10.1145/3607720.3607785.

[16] L. Lhotska, Why is ethics so important for artificial intelligence education?, in: 2025 34th Annual Conference of the European Association for Education in Electrical and Information Engineering (EAEEIE), IEEE, Piscataway, NJ, 2025, pp. 1–6. doi:10.1109/eaeeie65428.2025.11136589.

[17] M. Munarini, J. Brusseau, L. Angeli, Equitable AI audits: Evaluating the evaluators in today's world, in: Proceedings of the 17th International Conference on Theory and Practice of Electronic Governance, ACM, New York, NY, 2024, pp. 1–8. doi:10.1145/3680127.3680140.

[18] C. Todorova, G. Sharkov, H. Aldewereld, S. Leijnen, A. Dehghani, S. Marrone, C. Sansone, M. Lynch, J. Pugh, T. Singh, K. Mezei, P. Antal, P. Hanák, A. Barducci, F. Perez-Tellez, F. Gargiulo, The European AI tango: Balancing regulation innovation and competitiveness, in: Proceedings of the 2023 Conference on Human Centered Artificial Intelligence: Education and Practice, ACM, New York, NY, 2023, pp. 2–8. doi:10.1145/3633083.3633161.

[19] O. Aziukovskyi, V. Hnatushenko, V. Kashtan, A. Polyanska, A. Jamróz, Intelligent electricity load forecasting method using ARIMA-LSTM-Random Forest, Inżynieria Mineralna – Journal of the Polish Mineral Engineering Society 1 (1) (2025). doi:10.29227/IM-2025-01-18.

[20] I. A. Okwor, G. Hitch, S. Hakkim, S. Akbar, D. Sookhoo, J. Kainesie, Digital technologies impact on healthcare delivery: A systematic review of artificial intelligence (AI) and machine-learning (ML) adoption, challenges, and opportunities, AI 5 (2024) 1918–1941. doi:10.3390/ai5040095.

[21] S. R. Mishra, S. Dash, S. Padhy, P. Samuel, Legal aspects of operating IoMT applications in the fog computing, in: Advances in Science, Technology & Innovation, Springer, Cham, 2026, pp. 211–224. doi:10.1007/978-3-031-96265-3_15.

[22] F. Walke, L. Bennek, T. J. Winkler, Artificial intelligence explainability requirements of the AI Act and metrics for measuring compliance, in: Lecture Notes in Information Systems and Organisation, Springer, Cham, 2025, pp. 113–129. doi:10.1007/978-3-031-80122-8_8.

[23] E. Camilleri, Artificial Intelligence from Science Fiction to Reality, CRC Press, Boca Raton, FL, 2025. doi:10.1201/9781003653950.
[24] H. Billhardt, J.-A. Santos, A. Fernández, M. Moreno, S. Ossowski, J. A. Rodríguez, Streamlining advanced taxi assignment strategies based on legal analysis, Neurocomputing 483 (2022) 386–397. doi:10.1016/j.neucom.2021.10.085.