

Overcoming Resistance to Change: Artificial Intelligence in the Energy Sector

Jerome Lambert* 

Graduate School of Management, Saint-Petersburg State University, Russia

ABSTRACT

Background and Objectives: Artificial Intelligence (AI) promises productivity, safety, and sustainability gains in asset-intensive sectors; however, outcomes in the energy sector remain uneven. The sector's safety-critical operations, capital intensity, and stringent regulatory requirements make it a particularly demanding context for AI adoption, where technical performance alone is insufficient to ensure value. This study treats adoption as a socio-technical process rather than a tooling decision. It addresses three research questions: (RQ1) how workforce development and change leadership shape acceptance and sustained use; (RQ2) which organisational and governance conditions mitigate resistance and enable legitimate deployment; and (RQ3) under what conditions adoption yields operational reliability and environmental performance aligned with decarbonisation goals.

Methodology: A qualitative, multi-case design triangulated a semi-structured interview with a senior manager, Likert-scale surveys of mid-level managers and technical staff, and analysis of internal policies and strategy documents. Data were anonymised, thematically coded using a blended inductive–deductive approach, organised in a shared codebook, and synthesised across cases to map convergences and divergences in readiness, workforce development, and governance. Intercoder reliability was assessed, and disagreements were resolved through adjudication and iterative refinement of the codebook across cases. Triangulation maintained a transparent chain of evidence. Ethical safeguards included obtaining informed consent, maintaining confidentiality, and obtaining prior approval from the relevant institutional authorities.

Main Results: Three reinforcing levers shape adoption outcomes. First, broad-based capability building beyond specialist teams prevents benefits from concentrating in expert enclaves and reduces brittle scale. Second, communicative governance that couples transparency with contestability, through model cards, bias tests, validation reports, and explicit appeal rights, earns trust, curbs shadow workarounds, and improves safety culture. Third, a tight workflow fit that minimises

ARTICLE INFO

Article history:

Received 8 September 2025

Revised 18 March 2026

Accepted 20 March 2026

Keywords:

Artificial Intelligence,
Energy Sector,
Workforce Development,
Ethical Governance,
Sustainability

cognitive overhead at the decision point accelerates legitimate use and strengthens links to emissions monitoring and predictive-maintenance outcomes. Thin training coverage fosters anxiety about substitution and slows diffusion; structured upskilling and precise recourse mechanisms are associated with higher confidence, productivity, and clearer sustainability pathways.

Discussions: Algorithmic accuracy alone does not determine value; legitimacy and uptake hinge on people's and process readiness. The three levers translate literature on dynamic capabilities, AI readiness, and human responses to automation into operational guidance: invest in non-specialist literacy, institutionalise assurance and recourse, and engineer for workflow ergonomics in safety-critical contexts. Environmental gains materialise where oversight intensity, data quality, and targeted use cases align, indicating that governance quality conditions the conversion of adoption into credible emissions reductions. A responsible scale is pragmatic: build organisation-wide competence, communicate for legitimacy, and design for workflow fit.

Conclusions: Leaders should fund training coverage and design rather than headline hours, equip non-specialists to interpret model outputs, pair performance artefacts with participatory routines, and treat explainability as a usability requirement. Policymakers can reinforce these conditions by shifting from technology-neutral principles to auditable process standards that couple AI investment with reskilling and data-quality obligations. Future research should extend the design longitudinally and incorporate behavioural metrics to test causal links. The contribution is a field-tested playbook linking human capability, assurance, and workflow design to durable, auditable value in safety-critical energy contexts.

**Corresponding author*

E-mail address: j.lambert@gsom.spbu.ru

Introduction

Artificial Intelligence (AI) is rapidly diffusing across asset-intensive industries, with energy companies among the keenest adopters, since AI can optimise generation dispatch, transmission planning, and field operations, while also enabling emissions monitoring and predictive maintenance. Typical applications span forecasting (load, renewables, and prices), anomaly detection in rotating equipment, computer-vision-assisted inspections, and optimisation of maintenance windows within safety constraints. Emerging use extends to grid flexibility management, outage prediction under extreme weather, and coordination of

distributed energy resources, where milliseconds-level decisions and legacy operational-technology (OT) interfaces complicate deployment. Nevertheless, realised benefits vary markedly: ostensibly similar tools yield divergent organisational outcomes depending on complementary assets, managerial orientation, data quality, and the maturity of AI-specific capabilities needed to operate models at scale. This dispersion underscores that adoption is not merely a technical purchase, but a transformation of routines, decision-making rights, and accountability, often unfolding across heterogeneous asset classes and mixed vendor ecosystems.

Resistance to change is a principal reason for this variance. Employees and managers experience anxiety, mistrust, or identity threats when AI reshapes roles and decision rights; without credible communication and clear pathways for skill renewal, diffusion stalls even where the business case is strong. Mechanisms include perceived surveillance, loss of autonomy, role ambiguity, and uncertainty about error liability in safety-critical contexts, compounded by union expectations and stringent incident-reporting regimes. Recent evidence also suggests that stress associated with AI can trigger either adaptive learning or avoidance, depending on the trust and leadership context, thereby reinforcing the need to design augmentation rather than substitution into work systems and to stage reskilling alongside deployment (Zhou et al., 2024).

These dynamics are particularly acute in the energy sector, where safety-criticality, capital intensity, stringent incident-reporting regimes, union expectations, and the irreversibility of operational errors all heighten the perceived stakes of adopting AI. Unlike retail or professional services, energy operations involve high-consequence, real-time decisions where algorithmic failure can endanger lives, infrastructure, and the environment, making resistance both rational and deeply rooted.

The energy transition brings the global relevance of these challenges into sharp relief. As jurisdictions from the European Union to Southeast Asia accelerate decarbonisation mandates, AI is increasingly positioned as a key enabler of emissions monitoring, grid optimisation, and asset lifecycle management. Understanding and overcoming resistance to AI in energy is therefore not merely a firm-level concern but a precondition for meeting climate commitments at scale.

Sectoral conditions amplify these tensions. Energy operations are safety-critical, capital-intensive, and tightly regulated, making explainability, auditability, and lifecycle governance prerequisites for legitimacy and scale, rather than optional add-ons (Tchente et al., 2023). Public-sector-adjacent governance further shapes how utilities and state-owned enterprises specify, monitor, and improve AI, with uneven capabilities across jurisdictions in data stewardship, procurement, and oversight (Wirtz et al., 2022). Policy design also matters. Quasi-natural experiments suggest that targeted AI initiatives can reduce firms' carbon emissions. However, the effects are heterogeneous and firm-specific, reflecting baseline capabilities, regulatory exposure, and organisational incentives (Feng et al., 2024). Compliance obligations, model documentation, and audit trails, therefore, become operational requirements that influence vendor selection, update cadence, and ultimate scale.

Complementarities are the second source of outcome dispersion. Studies show that AI's value materialises when paired with a robust data infrastructure, integrated OT/IT architectures, and adjacent digital technologies such as IoT sensing and big data analytics. Where these complements are present, growth and efficiency effects are more potent because models can be monitored, versioned, and retrained reliably. From a capability perspective, organisations that

purposefully develop AI-specific resources and dynamic routines, data pipelines, MLOps, model risk controls, and cross-functional teaming experience greater innovation and performance, underscoring that "AI capability" is an organisational property rather than a tool purchase. At the same time, end-user explainability must shape acceptance and sustained use, necessitating human-centred interfaces and decision support aligned with operator workflows (Haque et al., 2022). Brownfield integration, cyber-security hardening, and data lineage controls further condition whether pilots generalise across fleets and geographies.

Sustainability outcomes likewise depend on managerial time horizons and institutional pressures. Energy firms with stronger ESG engagement and clear green strategies appear to leverage AI's innovation benefits more effectively. In contrast, short-term executive priorities can blunt or misdirect them. Translating adoption into measurable environmental performance requires governance that aligns incentives with carbon-reduction goals, specifies data and measurement protocols, and links projects to assets with material abatement potential (Wang et al., 2024).

Despite growing scholarly attention, the existing literature treats workforce development, governance design, and workflow integration as largely separate enablers of AI adoption. No prior study has examined how these three dimensions interact to shape adoption outcomes across energy firms operating under different ownership structures, market logics, and regulatory regimes. This gap is consequential because the energy sector's safety-criticality, asset intensity, and decarbonisation imperatives demand an integrated account that connects people readiness, institutional assurance, and operational ergonomics. This study addresses the gap by providing a multi-case analysis that bridges these streams and offers practical guidance for managers and policymakers seeking to scale AI responsibly.

Accordingly, this article addresses three research questions:

RQ1: How do workforce development practices and change leadership shape employee acceptance and sustained use of AI in energy operations?

RQ2: Which organisational and governance conditions mitigate resistance to AI and enable legitimate, safety-conscious deployment?

RQ3: Under what conditions does AI adoption translate into operational reliability and environmental performance aligned with decarbonisation goals? By analysing evidence from three organisations operating under different ownership structures and market logics, the study isolates how workforce development and change leadership influence acceptance and sustained use, which organisational and governance conditions mitigate resistance, and when adoption translates into operational and environmental performance. Interviews, surveys, and document data are triangulated to derive practical guidance for scaling AI responsibly in the energy sector, outlining actionable contributions for managers and policymakers. The Literature Review elaborates the theoretical background, and the Experimental Procedure details the qualitative, multi-company comparative design used to generate the findings.

Review of Literature

AI as a Socio-Technical System: Promise and Predispositions

Artificial Intelligence is now embedded across core organisational processes and public services, offering efficiency and prediction gains while also importing and amplifying the prejudices of the data, models, and people who build and operate it. These gains increasingly

encompass forecasting, optimisation, and decision support, with applications spreading from back-office analytics to frontline decision aids. Nevertheless, the same pipelines can encode historical exclusions, propagate labelling artefacts, and entrench skewed outcome distributions across populations as systems scale and are repurposed. Recent work grounded in a Brazilian context shows how citizens already recognise that "neutral" systems can reproduce societal discrimination through data quality and diversity problems, algorithmic inaccuracies, and developer/operator bias, conditions that have prompted regulatory initiatives such as Brazil's Bill No. 2338/23, which aims to secure fair, non-discriminatory AI. Parallel streams in industrial marketing frame "algorithmic bias management" as a capability spanning data, model, and deployment layers, highlighting that mitigation must be systemic rather than a patchwork of technical fixes (Akter et al., 2023). A complementary body of auditing literature suggests that principles such as fairness, transparency, and responsibility require operational definitions, measurable criteria, and assurance practices to be effective in real-world deployments, implying the need for governance that links metrics to remedial action and continuous monitoring (Laine et al., 2024).

Sources of Discrimination: Data, Models, and Organisational Choices

Across domains, evidence converges on three recurrent sources of discrimination: unrepresentative or low-quality data, brittle models that perpetuate historical disparities, and organisational choices that reflect the preferences of their creators. Data issues include sampling bias, missingness concentrated in vulnerable groups, and feedback effects, wherein automated decisions reshape future datasets, thereby worsening the imbalance over time. Model issues range from overfitting to spurious correlations and fairness–accuracy trade-offs that are rarely interrogated with domain-sensitive constraints or error-budget thinking. Organisational factors encompass target selection, thresholding, and incentive schemes that prioritise short-run efficiency over equity, often without explicit review. In recruitment and selection, for example, applicants often trust human assessors more than AI evaluators and report lower job acceptance intentions when AI screens them, associations explained by doubts that an AI can perceive "unique skills," which in turn raises fairness and legitimacy concerns if such systems are deployed at scale (Deriu et al., 2024). In customer-facing decision support, classic CRM analyses warn that AI-driven micro-targeting can harden "prioritisation" logics into service discrimination unless design constraints and oversight intervene (Libai et al., 2020). These patterns echo what citizens report in public forums and surveys: perceived accuracy, appeal mechanisms, transparency, and the absence of private-company involvement are more salient drivers of acceptance and perceived procedural fairness than generic promises of innovation (Horvath et al., 2023). Collectively, these findings underscore that technical excellence must be coupled with institutional safeguards to achieve durable legitimacy.

State Capacity and Public-Sector Adoption

The public sector's AI turn is therefore shaped as much by institutional capacity and governance design as by the algorithms themselves. Studies in India and Saudi Arabia identify barriers that range from data integrity and infrastructure gaps to ethics and cultural tensions, showing that the issues to which leaders direct their attention, and the resources they allocate, determine whether assimilation advances or stalls (Alshahrani et al., 2021). Beyond infrastructure, capability gaps in procurement, benchmarking, and benefit realisation often

impede scale, while fragmented accountability obscures who is responsible for model risk in high-stakes contexts. An integrative review of "AI-GOV" argues that robust government use must satisfy operational fitness (it works in context), epistemic alignment (its knowledge claims are warranted), and normative alignment (it accords with public values), a triad that helps diagnose why some high-performing tools still fail legitimacy tests (Straub et al., 2023). At the city level, comparative surveys suggest that ease of use influences attitudes. In contrast, perceived usefulness is closely tied to tangible civic outcomes, such as safety, crisis management, and cost reduction, with marked cultural differences across jurisdictions (Yigitcanlar et al., 2023). Simultaneously, citizens' privacy concerns remain foundational in any government use of AI for collective behaviour analysis, pressing for lawful design and explicit safeguards against surveillance creep, including data minimisation, purpose limitation, and proportionate oversight (Saura et al., 2022).

Trust, Explainability, and Contestability

Because acceptance hinges on trust, explainability and communicative practice are not secondary features but core controls. Cross-sector case studies demonstrate that trust in applied AI is "dichotomous": it inheres in both the technology and the innovating firm's communication, implying that technical assurance without democratic dialogue will underperform (Hengstler et al., 2016). In supply chains and B2B CRM, explainable AI enhances transparency and agility and can mitigate the "dark sides" of machine recommendations by making the reasons for recommendations inspectable and contestable (Gaczek et al., 2023). Extending these insights, organisations increasingly pair interpretability techniques with user-centred documentation (e.g., model cards) and structured appeal pathways, recognising that contestability and human-in-the-loop design are preconditions for sustained use in regulated settings. Effective practice, therefore, couples local, task-relevant explanations with calibrated uncertainty displays, incident playbooks, and training that builds shared mental models between developers and end users.

Organisational Readiness and Complementarities

Mitigation also depends on whether organisations are "AI-ready." Cross-country modelling isolates enablers, scientific output, internet infrastructure, and public consumption expenditures, which similarly shape readiness in developed and developing economies (Pramanik et al., 2024). Qualitative work with multinationals in South and Southeast Asia unpacks readiness into informational, environmental, infrastructural, process, data, and technological dimensions, underscoring that "being ready" is a socio-technical condition, not a procurement event (Tehrani et al., 2024). Within the firm, HRM research draws together technology, organisation, and people perspectives to show how post-adoption frictions, opacity, and system-level factors can derail promised benefits unless assimilation is managed as an iterative change process; when AI capabilities are embedded as dynamic capabilities and supported by a data-driven culture, environmental and performance outcomes improve. However, the effects vary across national cultures and contexts, warning against one-size governance (Wamba et al., 2023). These findings highlight the importance of MLOps, data governance maturity, and cross-functional teaming as complements to any individual model; in practice, readiness manifests as reproducible pipelines, clear ownership, service-level

objectives, and escalation paths for addressing ethical conflicts, all of which reduce integration risk.

Labour Implications and Human Responses

The labour side of adoption reveals ambivalence that matters for fairness and uptake. In hospitality, generative and predictive systems can threaten job security and spur turnover intentions, yet also stimulate voice behaviours and innovative work when organisational support and learning climates are strong; the same technologies can be perceived either as a challenge or hindrance based on locus of control and future orientation (Lin & He, 2024). Event-based studies further show that the 'AI event strength' of visible deployments shapes rumination and proactive adaptation, with autonomy buffers that organisations can build in by design. (Bai et al., 2025). These dynamics caution that bias-mitigation strategies are inseparable from the working conditions, training coverage, and participation mechanisms that shape who builds, maintains, and is evaluated by AI. Job redesign, transparent task allocation between humans and systems, and credible upskilling pathways, therefore, become central to both legitimacy and productivity, aligning workforce perceptions with augmentation narratives.

Sectoral Lessons from Healthcare

Healthcare studies offer a parallel lesson: attitudes are broadly positive, but conditional on factors such as privacy, explainability, validation, and role clarity. Concerns about empathy, professional autonomy, and accountability persist, even among experienced clinicians who work with AI (Vo et al., 2023). Implementation research with radiologists reveals that time delays, extra steps, and unstable performance are key barriers to adoption, reminding policy designers that fairness and safety necessitate attention to workflow, not just accuracy (Wenderott et al., 2024). Providers of diagnostic AI themselves report "stealth science" workarounds and digital ambidexterity to cope with the complexity of clinical systems, reinforcing the need for governance that aligns with institutional realities (Gillner, 2023). Together, these studies suggest that effectiveness hinges on socio-technical fit: validated performance must be matched by integration into existing routines, clear escalation when models err, and delineated liability.

From Principles to Practice: Operational Governance

On the governance side, synthesis efforts converge on the need to translate principles into structures, routines, and audits. A scoping review conceptualises "responsible AI governance" as a set of structural, relational, and procedural practices with identifiable antecedents and effects. In contrast, an organisational AI governance review calls for explicit oversight units, clearer process models, and context-aware regulation rather than technology-neutral abstractions (Papagiannidis et al., 2025). The ethics-based AI auditing literature, for its part, identifies stakeholders, developers, deployers, users, regulators, and the public, as well as the knowledge contributions audits produce, making practical the oft-stated demand for accountability (Laine et al., 2024). In operational terms, this translates into model inventories, change-control gates, tiered risk assessment, and independent review cycles that connect findings to remediation and learning, ensuring traceability across the model lifecycle.

Environmental and Societal Externalities

Finally, the societal externalities of AI, both environmental and social, are measurable and governable. Several quasi-experimental studies have linked AI capabilities to lower corporate greenwashing and improved environmental performance through mechanisms such as productivity gains, reduced information asymmetries, and enhanced green innovation. Importantly, oversight intensity and institutional context moderate these effects, suggesting that "AI as a climate tool" is contingent upon the quality of governance. At the macro level, country-panel work associates AI diffusion with reductions in carbon emissions via efficiency gains, an effect that is amplified when regulation is coherent (Cao et al., 2025). These benefits, however, require disciplined data generation, auditable measurement systems, and alignment between operational projects and strategic sustainability targets to avoid rebound effects. Taken together, the literature positions discriminatory bias not as an aberration but as an emergent property of socio-technical systems, shaped by data provenance, model design, organisational incentives, and institutional choices. It also offers a practical palette of responses, including lifecycle bias-management capabilities, participatory and transparent designs that foster citizen legitimacy, explainable tools and audits that render decisions transparent, and governance architectures that align operations with democratic values.

Method

Research Design

The study adopted a qualitative, multi-case design to investigate how AI is being embedded within the energy sector. This approach permitted an in-depth exploration of organisational practices, managerial orientations, and governance arrangements across capital-intensive, safety-critical, and heavily regulated contexts. Cases were selected purposively to maximise variation across ownership structures, energy mixes (including renewable, nuclear, and fossil-fuel operations), and levels of AI maturity, thereby enabling theoretical replication and cross-case contrast (Straub & al., 2023). Selection criteria, rationale, and screening memos were systematically recorded in a case study database to strengthen reliability and ensure transparency of case inclusion decisions. Whilst this design supports rich comparative analysis, the sample of three firms limits statistical generalisation. The findings should therefore be understood as analytically generalisable rather than representative of the sector as a whole.

Participants

The inquiry focused on three large energy organisations, each embodying different ownership logics and operating conditions. Evidence was collected through three complementary streams of investigation. First, a semi-structured interview lasting approximately 25 minutes was conducted with a senior manager, focusing on AI readiness, workforce reskilling strategies, governance of ethical risks, and alignment with sustainability (Wenderott & al., 2024). Second, structured surveys were administered to mid-level managers and technical staff across two organisations. These surveys probed attitudes toward AI adoption, adaptability, perceived substitution risks, and ethical concerns. Third, documentary sources were gathered, including AI roadmaps, training curricula, data governance guidelines, and ESG disclosures. Together, these sources yielded a triangulated dataset reflecting both individual-

level perceptions and institutional-level commitments. Table 1 summarises case characteristics and data collected.

Table 1

Case Characteristics and Data Summary

Organisation	Ownership Type	Energy Mix	AI Maturity	Data Collected
EDF	State-majority owned	Nuclear, renewables, fossil	Advanced (pilots + partial scale)	Semi-structured interview (1 senior manager, ~25 min); AI roadmap, training records, ESG reports
Lukoil	Publicly traded, private, state-influenced	Oil, gas, refining	Early–intermediate	Survey ($n = 10$); data governance policy, strategy documents
Gazprom	State-controlled	Gas, oil, and power generation	Advanced (structured rollout)	Survey ($n = 12$); data governance policy, strategy documents

Instruments and Their Validation

Surveys were designed using a five-point Likert scale (strongly disagree–strongly agree). Several reverse-coded items were included to reduce acquiescence bias, and optional free-text prompts were provided to provide contextual detail (Tehrani & al., 2024). Constructs were defined a priori and refined iteratively during pilot testing. "AI readiness" covered leadership sponsorship, data-infrastructure quality, and MLOps practices; "workforce development" included training coverage, role clarity, and human-in-the-loop responsibilities; "governance" encompassed assurance artefacts, bias checks, model cards, and appeal mechanisms; "sustainability alignment" indexed links between AI use cases and emissions reduction or maintenance efficiency. Instruments were piloted with a small convenience sample to assess clarity, cognitive burden, and interpretability; item-level feedback led to minor wording revisions before full deployment. Pilot responses showed acceptable internal consistency across constructs, though formal Cronbach's alpha is not reported, given the small, purposive samples and the primarily qualitative role of the survey data. Documentary sources were appraised for authenticity, completeness, and recency (Laine & al., 2024).

Data Collection and Handling

All interviews were audio-recorded with consent, and detailed notes were transcribed for analysis. Survey responses were collected digitally, anonymised, and stored on encrypted media. Personally identifiable information was removed, and illustrative quotations were lightly paraphrased to reduce re-identification risk. Internal documents were systematically catalogued, with metadata on provenance and date of issue, ensuring the corpus remained current and relevant. Careful handling procedures were maintained to preserve confidentiality while facilitating the integration of analytics across data types.

Data Analysis

Analysis proceeded through a thematic approach combining inductive and deductive coding (Vo et al., 2023). An initial codebook was derived from the literature and the case protocol, then iteratively refined through constant comparison. Coding was first applied within cases to identify emerging themes and subsequently extended across cases through the development of cross-case matrices. This procedure supported both pattern matching and explanation building (Alshahrani et al., 2021). A subset of transcripts and survey responses was dual-coded, and discrepancies were adjudicated through analytic memos. Decision logs documented code merges, splits, and refinements, thereby enhancing dependability and reliability. Triangulation across interviews, surveys, and documents, combined with an audit trail of analytic decisions, maintained a transparent chain of evidence. Descriptive statistics (means, modal responses, and ranges) were computed for each construct and reported alongside qualitative themes.

Trustworthiness and Rigour

Credibility was reinforced through several strategies. Negative-case analysis was undertaken to identify counter-patterns, and rival explanations were systematically considered. Reflexive attention was given to the sectoral constraints of safety-critical, regulated operations, which shape both data quality and interpretive framing. Detailed case descriptions supported Transferability, while confirmability was enhanced through systematic documentation of analytic steps. Triangulation across sources, roles, and data types increased robustness and reduced the likelihood of single-source bias.

Ethical Considerations

Ethical safeguards were integral throughout the study. Prior approval was obtained from an institutional review board. All participants received informed consent documentation outlining the study aims, the voluntary nature of participation, and their right to withdraw at any time. Confidentiality guarantees were upheld through secure, role-restricted storage of de-identified data. Data handling complied with the General Data Protection Regulation (GDPR) for EU-based participants and Federal Law No. 152-FZ on Personal Data for Russia-based participants; institutional data-protection protocols were followed at all sites. Sensitive issues, including perceived substitution risks and algorithmic governance, were broached using neutral, non-leading phrasing, and debriefings were offered to minimise participant distress. Paraphrasing of quotations further reduced the risk of re-identification while retaining conceptual integrity.

Findings and Discussion

Adoption outcomes diverged less along sub-sectoral lines than by the depth of people-and-process readiness accompanying technical rollout. Across cases, AI achieved traction when three conditions coalesced: (i) broad-based competence building that equips non-specialists to interpret and contest model outputs; (ii) communicative governance that renders recommendations inspectable and establishes credible challenge rights; and (iii) workflow-tight design that minimises cognitive switching costs at the point of use. In practice, competence-building involved role-relevant literacy, hands-on exercises, and shared glossaries, thereby reducing reliance on gatekeeping experts. Communicative governance required visible artefacts,

routine forums for questioning, and traceable resolutions. Workflow-tight design was based on latency budgets, alarm thresholds, and interface cues aligned with existing safety procedures. Where any one of these was thin, organisations reported brittle scale, shadow workarounds, or contested legitimacy, even when technical performance appeared strong on paper, because perceived risk, rather than test accuracy, dominated frontline behaviour.

EDF illustrates a "high-tech/low-diffusion" equilibrium. AI supports roughly 45% of operations with reported downtime reductions nearing 20%, yet only about 40% of employees receive AI-specific training. The consequence is the concentration of benefits within expert enclaves and operational dependency on a small cadre of "AI translators." As one survey respondent noted in the free-text field, "the tools work well for those of us who built them, but the field teams still revert to the old checklist because they do not understand what the model is telling them." In safety-critical contexts, this creates a two-speed system: local model improvements compound within specialised teams, but cross-functional assimilation slows where explainability and role clarity are prerequisites. Interviews and documents suggest that assurance artefacts (validation reports, bias checks, model cards) exist but circulate primarily within engineering, limiting broader sensemaking and narrowing opportunities for pre-deployment rehearsal. The trajectory signals latent potential: expanding foundational literacy for planners, operators, and maintenance staff, coupling it with routinised pre-deployment walk-throughs, structured post-incident reviews, and periodic "challenge clinics," would shorten the distance between model output and frontline judgement, reduce reversions to prior practice, and dilute key-person risk associated with translators.

Lukoil exemplifies demand-side frictions that stall diffusion even when the efficiency case is clear. Surveyed staff reported notable anxiety, with ~35% fearing job displacement and 55% feeling unprepared for AI-related tasks. Evidence of "threat appraisals" surfaced: automation was framed as substitution, and appeal mechanisms were opaque; operators reported defensive resistance and stronger intentions to bypass tools. One mid-level manager wrote: "We were told the system would 'support' our decisions, but nobody explained what happens when we disagree with it or whether overriding it counts against us." Conversely, when demonstrations highlighted human-in-the-loop guardrails and task-level autonomy, respondents described a willingness to experiment and to surface edge cases. This asymmetry highlights the importance of communication quality as a key operational lever, not merely an add-on. Without visible skill pathways (e.g., micro-credentials mapped to concrete role transitions), scenario-based drills that rehearse override use, and accessible recourse (documented protocols with no-penalty review), identical tools produce opposite behavioural responses: compliance without engagement versus engaged adoption that yields learning effects and safer escalation.

Gazprom sits at the opposite end of the readiness spectrum. Structured, inclusive training has reached over 60% of the workforce, featuring role-based curricula tailored for operators, analysts, and supervisors. As understood from the research: "We invested in role-specific curricula early, not because it was required, but because we saw that operators who understood the model's limits trusted it more and escalated more precisely." AI is integrated into emissions monitoring and efficiency initiatives. The reported gains in confidence, job satisfaction, and productivity are consistent with an organisation-level capability: reliable data pipelines, disciplined model operations, change roles embedded in line functions, and cross-functional teaming rather than a product purchase. Crucially, governance ties use cases to

strategy and measurement: dashboards track not only accuracy and latency but also override rates, time-to-decision, alert fatigue, and post-incident learning actions. This instrumentation furnishes 'evidence of fit' rather than mere assertion, lowers contestation costs, enables targeted retraining, and accelerates iteration in regulated environments where auditability is decisive for scale.

Table 2

Cross-Case Comparison Matrix

Dimension	EDF	Lukoil	Gazprom
Training Coverage	~40% of staff; concentrated in engineering teams	Limited; ~55% feel unprepared; no visible reskilling pathway	>60% of workforce; role-based curricula for operators, analysts, supervisors
Governance Artefacts	Model cards and validation reports exist, but circulate within engineering	Appeal mechanisms are opaque; there are no documented override protocols	Dashboards tracking accuracy, overrides, alert fatigue, tied to strategy
Workflow Fit	Partial; AI outputs not fully integrated into frontline routines	Weak; tools perceived as cognitive overhead competing with safety practices	Strong; co-designed alarm logic, escalation paths, uncertainty displays
Sustainability Alignment	Unrealised potential; strong technical base but limited diffusion	Cost-optimisation framing dominates; cultural resistance neutralises efficiency	AI linked to emissions monitoring, energy intensity; auditable KPIs
Overall Outcome	High-tech / low-diffusion equilibrium	Demand-side resistance stalls adoption	Broad-based, safety-conscious scale

Three cross-case mechanisms explain outcome variance. First, training coverage and design matter more than the number of hours. Programmes that build baseline literacy for non-specialists (terminology, uncertainty, data quality, model limits, and simple back-of-the-envelope checks) reduce the interpretive gap that otherwise forces decisions back into tacit heuristics and diminish reliance on scarce translators. Second, legitimacy is co-produced by performance and communication. Assurance artefacts travel when paired with participatory routines—pre-deployment simulations that surface edge cases, opt-out clauses for poorly modelled states, explicit appeal rights with tracked resolution times, and "model office" sessions where operators question developers and see changes implemented. Third, workflow fit is a fairness and safety issue as much as a productivity one. Where interfaces add steps, enlarge handoffs, or increase latency variance, operators perceive tools as cognitive overhead that

competes with established safety practices; where escalation paths, alarm logic, and uncertainty displays are co-designed, tools are experienced as safety assets that preserve professional discretion.

These people-centric levers facilitate the environmental pathway. When AI targets emissions monitoring, energy intensity, or predictive maintenance, benefits depend less on initial model accuracy and more on disciplined data generation, lifecycle monitoring, and alignment with oversight. This finding helps explain why similar tools produced different sustainability payoffs under distinct ownership and regulatory exposures: structured upskilling, combined with targeted use cases, yielded clearer links to environmental performance; strong technical foundations, without commensurate reskilling, signalled unrealised potential; cost-optimisation framings, absent anticipatory reskilling, invited cultural resistance that neutralised technical efficiency. Pairing operational KPIs with environmental indicators (avoided flaring events, leakage-detection precision/recall, maintenance-related energy savings, and time-to-mitigation) made claims auditable, anchored benefit-tracking in shared dashboards, and focused improvement cycles on verifiable reductions rather than narrative wins.

Theoretical Contribution

The findings extend two theoretical traditions. First, they operationalise the dynamic capabilities framework (Teece, 2007) in the context of safety-critical energy operations. The three levers, competence building, communicative governance, and workflow fit, correspond to the sensing, seizing, and transforming micro-foundations of dynamic capabilities. Competence building enables sensing by equipping non-specialists to detect model anomalies and contextual misfit. Communicative governance supports seizing by providing the legitimacy and trust infrastructure needed to commit resources. Workflow fit enables transformation by embedding AI into routines, so reconfigured practices become durable rather than episodic.

Second, the study enriches socio-technical systems theory (Trist & Bamforth, 1951; Baxter & Sommerville, 2011) by demonstrating that the joint optimisation of technical and social subsystems must extend beyond initial design to encompass lifecycle governance, participatory assurance, and iterative skill renewal. The cross-case evidence shows that partial optimisation, strong technical foundations without commensurate social scaffolding, or vice versa, produces brittle outcomes. The three-lever framework thus offers a mid-range operationalisation that connects abstract socio-technical principles to concrete managerial actions in regulated, high-consequence settings.

Managerially, the findings specify a sequence. Invest early in competence and operating routines, baseline literacy, override and appeal protocols, alarm-logic reviews, and post-incident learning loops, then scale model complexity and scope as adoption stabilises. Treat explainability as a usability requirement: at the decision point, operators need inputs, uncertainty bounds, salient factors, and plausible alternatives within their latency budget. Institutionalise artefacts (model cards, bias tests, validation reports) as living documents used in governance forums, and shift them from compliance archives to operational debate tools. Resource product-like roles embedded in operations to coordinate between developers and line functions, maintain a visible backlog of model improvements, and balance success measures that combine accuracy, adoption depth, safety incidents, and environmental impact.

For policy and oversight, attention should shift from technology-neutral principles to auditable process standards. The implications differ across regulatory regimes. In the European

Union, the AI Act's risk-tiered framework provides a foundation, but energy-specific implementing standards are needed to operationalise human oversight and documentation requirements for safety-critical applications. In the United States, the sectoral approach relies on existing agencies (FERC, NRC) to integrate AI governance into domain-specific regulation, creating potential gaps when cross-cutting issues, such as workforce readiness, fall between jurisdictions. In Asia, advisory frameworks (e.g., ASEAN's AI governance guides) offer flexibility but may lack enforcement teeth for high-consequence energy applications. These differences underscore the importance of context-aware governance rather than one-size-fits-all prescriptions.

Connecting these findings to broader ESG compliance frameworks, organisations that embed AI governance within their environmental, social, and governance reporting—documenting workforce reskilling, override analytics, and emissions-linked KPIs—are better positioned to satisfy investor scrutiny, regulatory audits, and the energy transition's demand for credible decarbonisation pathways.

Conclusion

Across large energy incumbents, AI performance is less a function of algorithms than of the socio-technical scaffolding surrounding them. The comparative evidence shows that outcomes hinge on three mutually reinforcing levers: broad-based capability building beyond specialist teams, communicative governance that earns trust through transparency and recourse, and tight workflow fit that reduces cognitive and temporal overhead at the point of use. When these conditions are present, adoption transitions from isolated pilots to a reliable, safety-conscious scale, and it is more likely to yield operational gains alongside measurable environmental benefits that extend beyond efficiency improvements. Conversely, where scaffolding is partial, skills are concentrated in expert enclaves, governance is treated as compliance paperwork, or interfaces burden operators, organisations encounter brittle scale, uneven legitimacy, and a reversion to pre-AI routines. Such failures often erode workforce confidence, deepen perceptions of unfairness, and weaken public trust, making recovery costlier than designing for readiness from the outset. In short, value is produced by orchestration across people, processes, and tools, not by model performance in isolation, and legitimacy becomes as important as technical reliability in sustaining momentum.

Managerially, the priority is to fund training coverage and design, not only training hours; equip non-specialists with data and model literacy; institutionalise assurance artefacts (validation reports, bias checks, model cards) alongside participatory routines (walk-throughs, challenge rights, explicit appeals); and treat explainability as a usability requirement in safety-critical contexts. Aligning AI programs with emissions monitoring and maintenance use cases, backed by disciplined data generation and lifecycle metrics, strengthens the link between deployment and sustainability performance. Boards and public owners can accelerate this alignment by requiring evidence of workforce readiness and auditability in capital approvals. Practical enablers include role-based curricula, published override protocols with no-penalty review, post-incident learning loops, and product-owner roles embedded in operations to coordinate model evolution with frontline needs. Progress should be tracked with mixed indicators that pair accuracy and latency with adoption, override rates, and safety or

environmental outcomes, thereby balancing technical precision with human legitimacy and organisational resilience.

Policy and governance implications follow directly. Sector regulators and public buyers can shift attention from technology-neutral principles to auditable process standards, including documentation of training coverage, human-in-the-loop controls, and appeal mechanisms. Reporting regimes that pair model performance with people and process indicators offer a clearer view of legitimacy, not only accuracy, and help policymakers identify early warning signs of misuse or erosion of trust. Where decarbonisation is a statutory objective, procurement and incentive schemes that couple AI spending with reskilling and data quality investments are more likely to produce persistent emissions reductions. Complementary measures, such as model documentation requirements, incident and near-miss reporting, proportionate stress-testing, and retention of audit logs, create verifiable accountability and reduce contestation costs, thereby fostering both social acceptance and regulatory compliance. Outcome-based regulation tied to independently verifiable environmental indicators can further align firm incentives with public goals, ensuring that AI contributes to sustainability transitions rather than exacerbating rebound effects or greenwashing risks.

The study's scope, comprising three large, asset-intensive firms, one interview, two survey sites, and internal documents, limits external generalizability and relies partly on self-reported assessments. Future research should extend the design in several specific directions: (a) longitudinal panel designs tracking adoption, workforce confidence, and safety outcomes over two-to-three-year windows; (b) incorporation of directly observed behavioural data, such as system interaction logs, override frequencies, and time-to-decision metrics, to complement self-reports; (c) quasi-experimental designs exploiting staggered rollouts within multi-site operators to sharpen causal attribution; (d) cross-industry replication (e.g., mining, petrochemicals, transport) to test whether the three-lever framework generalises beyond energy; and (e) formal psychometric validation of survey instruments with larger samples, including Cronbach's alpha and confirmatory factor analysis. Comparative studies across regions with different regulatory maturities and labour regimes could further illuminate whether the observed mechanisms are globally generalisable or contingent on institutional context.

Overall, the pathway to responsible scale in energy is pragmatic and actionable: build organisation-wide competence, communicate for legitimacy, and engineer for workflow fit. Firms that treat AI as an organisational capability, rather than a set of tools, are best positioned to convert adoption into resilient operations and credible progress on climate goals. Sustained advantage will accrue to organisations that institutionalise these routines, measure them transparently, and refine them continually as operating conditions and regulatory expectations evolve. In doing so, they not only future-proof operations but also contribute to a broader model of responsible innovation, where AI delivers measurable societal and environmental returns alongside efficiency and profitability.

Author Contributions

The single author of this paper has made the entirety of this article.

Declaration of the Use of Generative AI

During the preparation of this manuscript, the author used Grammarly. The tool was used solely to enhance grammatical correction of the text.

Ethics

This study received ethical approval from Saint-Petersburg State University. All procedures followed relevant ethical guidelines, and informed consent was obtained from all participants.

References

- Akter, S., Sultana, S., Mariani, M., Wamba, S. F., Spanaki, K., & Dwivedi, Y. K. (2023). Advancing algorithmic bias management capabilities in AI-driven marketing analytics research. *Industrial Marketing Management*, 114, 243–261. <https://doi.org/10.1016/j.indmarman.2023.08.013>
- Alshahrani, A., Dennehy, D., & Mäntymäki, M. (2021). An attention-based view of AI assimilation in public sector organisations: The case of Saudi Arabia. *Government Information Quarterly*, 39(4), 101617. <https://doi.org/10.1016/j.giq.2021.101617>
- Bai, J. Y., Huan, T. C. T., Leong, A. M. W., Luo, J. M., & Fan, D. X. (2025). Examining the influence of AI event strength on employee performance outcomes: Roles of AI rumination, AI-supported autonomy, and felt obligation for constructive change. *International Journal of Hospitality Management*, 126, 104111. <https://doi.org/10.1016/j.ijhm.2025.104111>
- Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, 23(1), 4–17. <https://doi.org/10.1016/j.intcom.2010.07.003>
- Cao, Q., Chi, C., & Shan, J. (2025). Can artificial intelligence technology reduce carbon emissions? A global perspective. *Energy Economics*, 143, 108285. <https://doi.org/10.1016/j.eneco.2025.108285>
- Deriu, V., Pozharliev, R., & De Angelis, M. (2024). How trust and attachment styles jointly shape job candidates' AI receptivity. *Journal of Business Research*, 179, 114717. <https://doi.org/10.1016/j.jbusres.2024.114717>
- Feng, L., Qi, J., & Zheng, Y. (2024). How can AI reduce carbon emissions? Insights from a quasi-natural experiment using generalised random forest. *Energy Economics*, 141, 108040. <https://doi.org/10.1016/j.eneco.2024.108040>
- Gaczek, P., Leszczyński, G., & Mouakher, A. (2023). Collaboration with machines in B2B marketing: Overcoming managers' aversion to AI-CRM with explainability. *Industrial Marketing Management*, 115, 127–142. <https://doi.org/10.1016/j.indmarman.2023.09.007>
- Gillner, S. (2023). We're implementing AI now, so why not ask us what to do?-How AI providers perceive and navigate the spread of diagnostic AI in complex healthcare systems. *Social Science & Medicine*, 340, 116442. <https://doi.org/10.1016/j.socscimed.2023.116442>
- Haque, A. B., Islam, A. N., & Mikalef, P. (2022). Explainable artificial intelligence (XAI) from a user perspective: A synthesis of prior literature and problematising avenues for future research. *Technological Forecasting and Social Change*, 186, 122120. <https://doi.org/10.1016/j.techfore.2022.122120>
- Hengstler, M., Enkel, E., & Duelli, S. (2016). Applied artificial Intelligence and trust—The case of autonomous vehicles and medical assistance devices. *Technological Forecasting and Social Change*, 105, 105–120. <https://doi.org/10.1016/j.techfore.2015.12.014>
- Horvath, L., James, O., Banducci, S., & Beduschi, A. (2023). Citizens' acceptance of artificial intelligence in public services: Evidence from a conjoint experiment about processing

- permit applications. *Government Information Quarterly*, 40(4), 101876.
<https://doi.org/10.1016/j.giq.2023.101876>
- Laine, J., Minkinen, M., & Mäntymäki, M. (2024). Ethics-based AI auditing: A systematic literature review on conceptualisations of ethical principles and knowledge contributions to stakeholders. *Information & Management*, 61(5), 103969.
<https://doi.org/10.1016/j.im.2024.103969>
- Libai, B., Bart, Y., Gensler, S., Hofacker, C. F., Kaplan, A., Köttlerheinrich, K., & Kroll, E. B. (2020). Brave new world? On AI and the management of customer relationships. *Journal of Interactive Marketing*, 51(1), 44–56. <https://doi.org/10.1016/j.intmar.2020.04.002>
- Lin, Q., & He, L. (2024). Does artificial intelligence (AI) awareness affect employees in giving a voice to their organisation? A cross-level model. *International Journal of Hospitality Management*, 123, 103947. <https://doi.org/10.1016/j.ijhm.2024.103947>
- Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems*, 34(2), 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
- Pramanik, P., Jana, R. K., & Ghosh, I. (2024). AI readiness enablers in developed and developing economies: Findings from the XGBoost regression and explainable AI framework. *Technological Forecasting and Social Change*, 205, 123482.
<https://doi.org/10.1016/j.techfore.2024.123482>
- Saura, J. R., Ribeiro-Soriano, D., & Palacios-Marqués, D. (2022). Assessing behavioral data science privacy issues in government artificial intelligence deployment. *Government Information Quarterly*, 39(4), 101679. <https://doi.org/10.1016/j.giq.2022.101679>
- Straub, V. J., Morgan, D., Bright, J., & Margetts, H. (2023). Artificial intelligence in government: Concepts, standards, and a unified framework. *Government Information Quarterly*, 40(4), 101881. <https://doi.org/10.1016/j.giq.2023.101881>
- Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13), 1319–1350.
<https://doi.org/10.1002/smj.640>
- Tchuente, D., Lonlac, J., & Kamsu-Foguem, B. (2023). A methodological and theoretical framework for implementing explainable artificial Intelligence (XAI) in business applications. *Computers in Industry*, 155, 104044.
<https://doi.org/10.1016/j.compind.2023.104044>
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3–38.
<https://doi.org/10.1177/001872675100400101>
- Tehrani, A. N., Ray, S., Roy, S. K., Gruner, R. L., & Appio, F. P. (2024). Decoding AI readiness: An in-depth analysis of key dimensions in multinational corporations. *Technovation*, 131, 102948. <https://doi.org/10.1016/j.technovation.2023.102948>
- Vo, V., Chen, G., Aquino, Y. S. J., Carter, S. M., Do, Q. N., & Woode, M. E. (2023). Multi-stakeholder preferences for the use of artificial Intelligence in healthcare: A systematic review and thematic analysis. *Social Science & Medicine*, 338, 116357.
<https://doi.org/10.1016/j.socscimed.2023.116357>
- Wamba, S. F., Queiroz, M. M., & Trinchera, L. (2023). The role of artificial intelligence-enabled dynamic capability on environmental performance: The mediation effect of a data-driven

- culture in France and the USA. *International Journal of Production Economics*, 268, 109131. <https://doi.org/10.1016/j.ijpe.2023.109131>
- Wang, Z., Zhang, T., Ren, X., & Shi, Y. (2024). AI adoption rate and corporate green innovation efficiency: Evidence from Chinese energy companies. *Energy Economics*, 132, 107499. <https://doi.org/10.1016/j.eneco.2024.107499>
- Wenderott, K., Krups, J., Luetkens, J. A., & Weigl, M. (2024). Radiologists' perspectives on the workflow integration of an artificial intelligence-based computer-aided detection system: A qualitative study. *Applied Ergonomics*, 117, 104243. <https://doi.org/10.1016/j.apergo.2024.104243>
- Wirtz, B. W., Weyerer, J. C., & Kehl, I. (2022). Governance of artificial Intelligence: A risk and guideline-based integrative framework. *Government Information Quarterly*, 39(4), 101685. <https://doi.org/10.1016/j.giq.2022.101685>
- Yigitcanlar, T., Li, R. Y. M., Beeramoole, P. B., & Paz, A. (2023). Artificial intelligence in local government services: Public perceptions from Australia and Hong Kong. *Government Information Quarterly*, 40(3), 101833. <https://doi.org/10.1016/j.giq.2023.101833>
- Zhou, Q., Chen, K., & Cheng, S. (2024). Bringing employee learning to AI stress research: A moderated mediation model. *Technological Forecasting and Social Change*, 209, 123773. <https://doi.org/10.1016/j.techfore.2024.123773>