

ORIGINAL RESEARCH

Open access

Corporate Governance in Algorithmically Mediated Organizations: Addressing Accountability, Transparency, and Strategic Oversight

Peter Novak^{1*}, Jana Svoboda¹

Abstract

The rapid integration of artificial intelligence and algorithmic systems into core organizational processes has transformed decision-making, yet it has simultaneously exposed critical deficiencies in traditional corporate governance mechanisms. Algorithmically mediated organizations now confront unique challenges in maintaining accountability for opaque automated decisions, ensuring transparency in high-stakes outcomes, and exercising strategic oversight amid rapid technological evolution. This conceptual manuscript synthesizes contemporary scholarship to map these tensions and introduces the SOAR Framework—Strategic Oversight for Algorithmic Responsibility—as a novel multi-layered governance architecture. Developed through a systematic review of peer-reviewed sources, the framework comprises six interdependent layers: the algorithmic decision core, transparency and explainability systems, accountability assignment protocols, strategic oversight bodies, risk and compliance shields, and adaptive feedback loops. By embedding human-AI hybrid controls and continuous audit mechanisms, SOAR enables organizations to align algorithmic mediation with ethical, legal, and strategic imperatives. The model addresses pressing gaps identified in existing literature, including the diffusion of responsibility in AI-driven environments and the insufficiency of conventional board-level oversight. Contributions to digital business and management studies include a practical blueprint for implementation and a conceptual foundation for future empirical testing. Ultimately, the SOAR Framework equips corporate leaders to govern algorithmic systems responsibly while preserving competitive advantage in digitally transformed enterprises.

Keywords Algorithmic governance, Corporate governance, Accountability, Transparency, Strategic oversight, AI-mediated organizations

*Correspondence:

Peter Novak
petr.novak@cvut.cz

¹ Department of Digital Economics and Innovation Systems, Charles University, Prague, Czech Republic

Introduction

Contemporary organizations increasingly delegate operational and strategic decisions to algorithmic systems, creating what scholars term algorithmically mediated organizations. From talent acquisition platforms that screen candidates without human intervention to dynamic pricing engines that set market strategies in real time, algorithms now permeate every functional domain. This shift promises unprecedented efficiency and data-driven precision [1-9];

however, it simultaneously erodes foundational pillars of corporate governance: accountability, transparency, and strategic oversight [8, 10-13].

Traditional governance models, rooted in hierarchical human oversight and periodic board reviews, prove inadequate when decision logic resides within black-box neural networks or proprietary reinforcement-learning agents [3, 14, 15]. Accountability becomes diffused when an algorithmic recommendation leads to financial loss or

ethical breach [6], raising questions about whether liability rests with developers, data scientists, deploying managers, or the board itself [11, 16–21]. Transparency is equally compromised [4]; stakeholders—investors, regulators, employees—cannot readily interrogate the rationale behind automated outcomes [14], undermining trust and regulatory compliance [18]. Strategic oversight, once exercised through deliberate deliberation [5], risks being reduced to post-hoc validation of machine-generated choices [16], limiting leadership's capacity to steer long-term organizational direction [22, 23].

These challenges are not theoretical. High-profile cases of algorithmic bias in hiring, discriminatory lending decisions, and autonomous pricing collusion have triggered regulatory scrutiny and reputational damage [7], demonstrating that governance failures carry material consequences [22, 24–27]. Moreover, the accelerating pace of AI advancement—exemplified by generative models and multi-agent systems—further widens the gap between technological capability and governance readiness [2, 19, 24].

Existing scholarship has documented these tensions across multiple streams. Research on AI governance highlights the need for organizational-level structures that extend beyond technical ethics guidelines [1, 5]. Studies of algorithmic management underscore the reconfiguration of power relations [10] and the emergence of new accountability voids [12, 25]. Corporate governance scholars, meanwhile, call for boards to develop AI-specific competencies and oversight protocols [8, 18, 21]. Yet a coherent, integrated framework that simultaneously addresses accountability, transparency, and strategic oversight within a unified architecture remains conspicuously absent [3].

This manuscript fills that void by proposing the SOAR Framework. Grounded in a synthesis of peer-reviewed publications [1], SOAR offers a multi-layered governance model specifically tailored to algorithmically mediated organizations [5]. The framework's six components create bidirectional flows between algorithmic operations and human strategic control [8], ensuring decisions remain explainable, attributable, and aligned with organizational purpose [12].

The paper proceeds as follows. The next section synthesizes the conceptual foundations drawn from AI governance, algorithmic accountability, and digital corporate strategy literatures. The subsequent section

introduces the SOAR Framework, detailing its architecture, components, and interrelationships, accompanied by a textual description of the governing visual model. By advancing this conceptual architecture, the manuscript equips scholars and practitioners with both a theoretical lens and a practical roadmap for responsible algorithmic mediation. In doing so, it contributes to the emerging discipline of digital business and management studies at a critical inflection point in organizational evolution.

Literature Synthesis and Conceptual Foundations

The literature on corporate governance in algorithmically mediated organizations has coalesced around three interlocking themes: the reconfiguration of decision rights, the imperative for explainability, and the redesign of oversight mechanisms [1]. Collectively, these streams reveal both the transformative potential of algorithmic systems and the profound governance deficits they engender [3].

Early contributions established that algorithmic mediation fundamentally alters organizational design [9]. Subsequent empirical and conceptual work formalized the concept of algorithmic management [10], showing that algorithms now perform coordination, evaluation, and direction functions traditionally reserved for human managers [12]. An extended insight argued that managing artificial intelligence requires new socio-technical arrangements that transcend conventional IT governance [17]. These studies collectively underscore a core proposition: algorithmic systems do not merely automate tasks [1]; they reallocate decision authority [3], necessitating governance innovation [5].

A parallel stream has scrutinized accountability deficits [3]. One examination revealed that reputational concerns often drive engagement strategies rather than structural reforms [6]. Further research highlighted the regulatory challenges of assigning responsibility when decision-making is distributed across human and machine actors [11, 13]. More recent work has demonstrated that fiduciary obligations must evolve to encompass oversight of automated decision architectures [27].

Transparency and explainability emerge as equally critical yet elusive requirements. One conceptualization frames transparency in AI as a multi-dimensional construct encompassing technical interpretability, procedural

openness, and stakeholder accessibility [15]. Research linking data governance to trustworthy AI shows that organizational data stewardship directly influences the feasibility of explainable outcomes [14]. These insights have been extended into wellbeing and social-responsibility domains [4], arguing that transparency is foundational to ethical algorithmic deployment [7]. Within strategic management literature, illustrations demonstrate how CIOs and boards can leverage explainability to maintain strategic control over AI-augmented processes [18, 20].

Corporate governance scholars have responded by advocating board-level AI competencies [8], coining the term “artificial governance” and positing that boards must evolve into hybrid human-AI oversight bodies [21]. Subsequent organizational AI governance frameworks emphasize alignment between technical capabilities and corporate strategy [1, 5]. Mapping knowledge gaps across AI governance literatures has identified the absence of integrated multi-layer models as a primary shortfall [8]. Complementary contributions from strategic information systems research underscore the lifecycle risks of algorithmic systems [25] and the necessity of continuous audit mechanisms [26].

Risk, compliance, and ethical considerations complete the foundational picture. Research demonstrates that AI-driven business model innovation introduces novel competitive and regulatory risks that traditional risk committees are ill-equipped to handle [16, 19, 23]. A theory of situated AI emphasizes that governance must be contextually embedded to preserve competitive advantage [24]. Regulatory perspectives further reinforce the need for company-law-compatible governance structures [2].

Orchestrating Responsible Algorithmic Stewardship: Introducing the SOAR Framework

To translate the synthesized conceptual foundations into organizational practice, this manuscript advances the SOAR Framework—Strategic Oversight for Algorithmic Responsibility [1]. SOAR is deliberately designed as a multi-layered architecture that creates bidirectional linkages between algorithmic operations and human strategic control [3], thereby operationalizing accountability,

transparency, and oversight within a single coherent model [5].

The framework comprises six interdependent layers. The foundational Algorithmic Decision Layer houses the core AI and machine learning systems that execute operational decisions [12, 17, 20]. Above it sits the Transparency and Explainability Systems layer, which mandates technical mechanisms—such as local interpretable model-agnostic explanations (LIME), SHAP values, and counterfactual reasoning—to render algorithmic outputs intelligible to stakeholders [4, 14, 15]. The third layer, Accountability Structures, assigns clear responsibility through hybrid human-AI protocols [3], role-based liability mapping [6], and automated audit trails that link decisions to accountable individuals or committees [11, 27].

The fourth layer, Strategic Oversight Mechanisms, elevates governance to the board and C-suite level via dedicated AI oversight committees that review high-impact decisions [1] and set algorithmic guardrails [8, 18, 21]. The fifth layer, Risk Management and Compliance Shield, integrates regulatory compliance checks [7], ethical impact assessments [13], and scenario-based stress testing to mitigate systemic risks [16, 23]. Finally, the Adaptive Feedback and Audit Loops layer closes the system through continuous monitoring, performance dashboards, and iterative refinement protocols that feed insights back into lower layers [5, 10, 25].

These layers do not operate in isolation. Vertical flows ensure that strategic directives cascade downward [8] while performance data and exception alerts flow upward [12]. Horizontal integration occurs through standardized interfaces that allow each layer to interrogate and influence adjacent layers in real time [14]. The architecture thereby prevents the fragmentation observed in current practice [9, 24, 26].

Figure 1 illustrates the SOAR framework as a six-layer governance architecture in which algorithmic decision systems are progressively translated into explainable, accountable, strategically supervised, risk-bounded, and continuously audited organizational action [1, 3, 5].




Figure 1. Governance Architecture for Algorithmically Mediated Organizations.

By embedding accountability at every layer, enforcing transparency as a design requirement, and institutionalizing strategic oversight as a continuous process, the SOAR Framework provides organizations with both a diagnostic tool and an implementation roadmap. Its dynamic, feedback-rich structure directly addresses the governance voids documented across the reviewed literature, offering a scalable solution for the algorithmically mediated enterprise.

Table 1 clarifies that each SOAR layer addresses a distinct governance question and operates through a different organizational mechanism, thereby demonstrating that responsible algorithmic governance requires functional translation rather than a single oversight intervention.

Table 1. Governance translation across the SOAR layers: functional role, primary governance question, organizational mechanism, and failure consequence

SOAR layer	Core governance function	Primary governance question addressed	Representative organizational mechanisms
Algorithmic decision layer	Executes automated or AI-assisted operational decisions	What decisions are being delegated to computational systems, and with what degree of autonomy?	Predictive models, recommendation systems, automated classifiers, reinforcement-based decision engines

Transparency and explainability systems	Converts technical outputs into interpretable and reviewable decision rationales	Can relevant stakeholders understand how and why the system produced a consequential outcome?	Model cards, decision logs, SHAP/LIME-based explanations, counterfactual outputs, traceability protocols
Accountability structures	Assigns ownership and responsibility across human-machine interfaces	Who is answerable for algorithmic outcomes, interventions, exceptions, and harms?	Role-based liability mapping, approval checkpoints, human-in-the-loop escalation rules, and decision provenance records
Strategic oversight mechanisms	Embeds algorithmic governance within executive and board-level control	How are algorithmic systems aligned with long-term strategy, fiduciary duty, and acceptable governance boundaries?	AI oversight committees, board review protocols, executive dashboards, delegated authority matrices, and policy guardrails
Risk management and compliance shield	Identifies, tests, and constrains legal, ethical, and systemic vulnerabilities	What harms, regulatory breaches, or systemic risks could emerge from algorithmic deployment under varied conditions?	Ethical impact assessments, compliance reviews, scenario stress testing, red team exercises, and control validation routines
Adaptive feedback and audit loops	Recalibrates systems and governance	How does the organization detect, learn	Continuous monitoring, periodic au

	structures through ongoing learning	from, and respond to performance deviations, incidents, and environmental change?	incident review boards, retraining protocols, and governance redesign triggers
--	-------------------------------------	---	--

Cultivating Algorithmic Accountability: Operationalizing the SOAR Layers in Practice

The SOAR Framework is not merely a conceptual architecture; it serves as a comprehensive, actionable governance blueprint that organizations can embed into everyday managerial routines, board-level decision protocols, and long-term technology strategies. Its central value lies in translating abstract governance principles into operational mechanisms that directly shape how algorithmic systems are designed, deployed, monitored, and recalibrated over time. By activating all six layers in an integrated manner, organizations can systematically close the governance gaps that currently expose algorithmically mediated firms to ethical failures, regulatory breaches, and strategic vulnerabilities [1, 8, 21]. Implementation, however, cannot occur through an abrupt structural overhaul. Instead, it requires a phased integration approach that aligns with existing governance infrastructures while progressively introducing hybrid controls to manage the speed, scale, and opacity of algorithmic decision-making. Through this gradual embedding, SOAR transforms governance from a periodic, compliance-driven function into a continuous, reflexive, and data-informed organizational capability.

Aligning legacy governance structures with algorithmic realities

Legacy corporate governance systems, traditionally anchored in periodic reporting cycles, hierarchical accountability, and functionally siloed compliance units, are increasingly misaligned with the temporal dynamics and epistemic opacity of algorithmic systems [18, 27]. Whereas boards historically relied on quarterly reviews and retrospective audits, algorithmic decision processes unfold in real time and often produce outcomes that resist straightforward ex post evaluation. This mismatch

necessitates a fundamental shift from episodic oversight to continuous and embedded governance. Within the SOAR framework, this shift is operationalized through the establishment of Algorithmic Stewardship Committees at the board level, structurally analogous to audit or risk committees but specifically mandated to oversee algorithmic systems as fiduciary concerns. These committees act as institutional anchors for the Strategic Oversight Mechanisms layer, receiving continuous data streams from the Adaptive Feedback and Audit Loops layer and translating them into strategic directives and corrective interventions.

In practice, this arrangement enables bidirectional accountability flows that are essential for governing complex sociotechnical systems. Strategic priorities and ethical thresholds are transmitted downward into operational layers, shaping model design, deployment parameters, and acceptable risk boundaries. Simultaneously, anomalies, performance deviations, and ethical concerns are escalated upward in near real time, ensuring that governance remains responsive rather than reactive. This dynamic flow directly mitigates the diffusion of responsibility documented in prior research on algorithmic decision-making environments, where accountability often becomes fragmented across technical and managerial domains [3, 6, 11]. By embedding algorithmic accountability into formal governance charters and board-level mandates, organizations institutionalize oversight that aligns with both fiduciary expectations and emerging regulatory pressures.

Developing hybrid human–machine oversight competencies

The effectiveness of any governance structure ultimately depends on the competencies of those who operate within it. A persistent challenge in algorithmic governance is the gap between technological complexity and managerial understanding, which often leads to overreliance on technical specialists or to uncritical acceptance of algorithmic outputs [8, 2, 18]. This competency gap manifests across multiple levels of the organization: directors lack the technical fluency to scrutinize AI-driven strategies, executives struggle to evaluate algorithmic performance metrics, and functional managers defer uncritically to system recommendations without adequate risk assessment [21, 27]. The consequences of this asymmetry are well documented, ranging from strategic

blind spots and regulatory non-compliance to reputational damage arising from undetected algorithmic failures [6, 13, 17].

The SOAR Framework addresses this challenge by foregrounding the development of hybrid human–machine oversight competencies as a core operational requirement. Rather than treating technical literacy as an optional skill, organizations are encouraged to institutionalize structured learning environments—such as cross-functional AI governance academies—that bring together directors, executives, legal experts, and data scientists in sustained engagement. These academies serve multiple functions: they provide foundational education on algorithmic architectures, facilitate scenario-based training on governance dilemmas, and create forums for cross-disciplinary dialogue that break down silos between technical and non-technical stakeholders [5, 20]. Such initiatives move beyond episodic training to embed continuous learning into the fabric of organizational governance, recognizing that algorithmic systems evolve rapidly and that competency development must be equally dynamic [25].

These initiatives serve to cultivate interpretive capacity, enabling decision-makers to understand not only how algorithmic systems function but also where their limitations and risks lie. Interpretive capacity encompasses the ability to question model assumptions, recognize when algorithmic outputs deviate from expected parameters, and identify conditions under which automated decisions may produce unintended or harmful outcomes [15, 24]. This capability is particularly critical in high-stakes domains such as hiring, credit allocation, healthcare, and criminal justice, where algorithmic errors can have profound consequences for individuals and expose organizations to significant liability [3, 11]. By fostering interpretive capacity, organizations move from passive reliance on algorithmic outputs to active governance that interrogates and validates system behavior.

At the same time, these initiatives foster ethical sensitivity by highlighting issues such as bias, discrimination, and unintended societal consequences, thereby reinforcing the fiduciary dimensions of algorithmic oversight. Ethical sensitivity involves recognizing when algorithmic systems implicate fundamental values such as fairness, privacy, autonomy, and non-discrimination [4, 7]. Through case-based learning and structured deliberation, decision-makers develop the capacity to identify ethical trade-offs

embedded in algorithmic design choices—for example, balancing predictive accuracy against explainability, or optimizing for efficiency at the expense of equity [14, 16]. This ethical orientation is not ancillary to governance but central to it, as algorithmic failures increasingly trigger regulatory sanctions, shareholder activism, and consumer backlash [2, 19].

The result is the emergence of hybrid accountability structures in which human judgment and machine intelligence operate in a complementary rather than substitutive relationship, consistent with the logic of the framework's third layer [27, 21]. Hybrid accountability structures distribute responsibility across human and algorithmic agents, leveraging the strengths of each while compensating for their respective vulnerabilities. Machines contribute speed, scalability, and pattern recognition; humans contribute contextual judgment, moral reasoning, and accountability for outcomes [10, 12]. This complementary arrangement resists the common pitfall of either abdicating responsibility to algorithms or unnecessarily constraining their potential through excessive human intervention [8, 22]. Instead, it establishes clear protocols for when and how human oversight should be exercised, including thresholds for escalation, override rights, and post-hoc review mechanisms [26].

This capability-building directly addresses the knowledge asymmetries identified in the literature, equipping leaders to interrogate algorithmic outputs, challenge underlying assumptions, and contextualize results within broader strategic and ethical considerations [5]. Leaders equipped with these competencies are better positioned to ask critical questions, such as: What data were used to train this model? What assumptions underpin its design? What are the known failure modes? How are errors detected and remediated? Who is accountable when things go wrong? [1, 23]. The capacity to pose such questions—and to evaluate the answers critically—transforms board-level oversight from a ceremonial function into an active and informed governance practice [9, 18].

The practical implications of this competency-focused approach extend across the six interdependent layers of the SOAR Framework. At the algorithmic decision core, interpretive capacity informs the design choices that shape system behavior [12]. Within transparency and explainability systems, ethical sensitivity guides decisions about what to disclose and to whom [15]. In accountability assignment protocols, hybrid structures clarify the

distribution of responsibility across human and machine actors [3, 11]. Strategic oversight bodies benefit from the cognitive diversity introduced by cross-functional training, reducing groupthink and enhancing scrutiny of technical decisions [8, 20]. Risk and compliance shields are strengthened when leaders can anticipate and mitigate algorithmic risks before they materialize [16, 19]. Finally, adaptive feedback loops are sustained by a culture that values learning from both successes and failures, enabling continuous refinement of governance practices [25, 26].

Deploying transparency mechanisms for stakeholder trust

Transparency within algorithmic systems is frequently treated as a secondary concern, addressed only after systems have been deployed and scrutinized. The SOAR Framework reconceptualizes transparency as an integral design principle that must be embedded directly within algorithmic infrastructures. The second layer operationalizes this principle by requiring the integration of explainability mechanisms into all high-stakes decision systems, thereby ensuring that algorithmic outputs are interpretable to both internal stakeholders and external audiences [4, 14, 15]. In practice, this involves deploying techniques such as SHAP values, counterfactual explanations, and natural-language interpretive interfaces, which collectively translate complex model behavior into accessible forms of reasoning.

However, transparency extends beyond internal interpretability. Organizations are also expected to engage in structured external disclosure by publishing Algorithmic Transparency Reports. These reports provide stakeholders with insight into how algorithmic systems influence critical organizational processes, including hiring, pricing, and credit allocation, while maintaining appropriate safeguards for proprietary information. By disclosing decision-making logic, governance structures, and documented incidents, firms signal their commitment to responsible algorithmic conduct and align with evolving regulatory expectations [2, 7, 13, 20]. Transparency thus serves as both a mechanism for accountability and a strategic resource for building trust, reducing reputational risk, and strengthening organizational legitimacy in environments marked by increasing scrutiny.

Fortifying risk shields against systemic vulnerabilities

The introduction of algorithmic systems fundamentally reshapes the nature of organizational risk, giving rise to complex and often emergent vulnerabilities that traditional compliance frameworks are ill-equipped to address. The Risk Management and Compliance Shield layer responds to this challenge by shifting the focus from reactive compliance to proactive resilience. Rather than relying solely on ex post audits, organizations embed ethical and risk assessments directly within the lifecycle of algorithmic systems, ensuring that potential harms are identified and mitigated before they materialize [13, 16].

This proactive orientation is operationalized through the integration of automated ethical impact assessments into development pipelines and the use of scenario-based stress testing to simulate extreme or adversarial conditions. These mechanisms enable organizations to anticipate how algorithmic systems might behave under conditions of data drift, strategic manipulation, or systemic interaction effects. Complementing these efforts, red-team exercises introduce an additional layer of scrutiny by deliberately probing systems, thereby uncovering hidden vulnerabilities and governance blind spots [23, 24].

Leveraging feedback loops for continuous governance adaptation

The outermost layer of the SOAR Framework completes the governance architecture by institutionalizing continuous adaptation. Algorithmic systems are inherently dynamic, evolving in response to changing data environments, user behaviors, and external conditions. Static governance mechanisms are therefore insufficient, as they fail to account for the ongoing recalibration required to maintain alignment between system performance and organizational values. The Adaptive Feedback and Audit Loops layer addresses this limitation by embedding continuous monitoring and learning processes into the core of governance operations.

Through regular algorithmic performance audits, organizations generate real-time insights into both technical and ethical dimensions of system behavior [5, 10, 25]. These insights are not merely diagnostic but are actively reintegrated into the algorithmic decision layer, enabling iterative refinement, recalibration, or, where necessary, decommissioning of systems that no longer meet acceptable standards. This closed-loop design ensures that governance evolves in parallel with technological change,

thereby preventing the accumulation of unmanaged risks and the persistence of outdated controls.

Table 2 demonstrates how the SOAR Framework can be used diagnostically by linking recurring governance failure modes in algorithmically mediated organizations to the specific control layers and remediation interventions required.

Table 2. Governance failure modes in algorithmically mediated organizations and corresponding SOAR control responses

Governance failure mode	Organizational manifestation	Underlying governance deficit	SOA layer(s) primarily activated		decisions ex post		
Opaque high-stakes decisions	Hiring, lending, pricing, or promotion outcomes cannot be meaningfully explained to affected stakeholders	Missing or inadequate interpretability infrastructure	Transparency and Explainability System; Strategic Oversight Mechanism	Model drift and governance obsolescence	Regulatory and ethical checks occur only after incidents or external complaints	Compliance is not integrated into the system lifecycle	Risk Management and Compliance Shield; Transparency and Explainability System
Responsibility diffusion after harm	Executives, developers, vendors, and managers each deny ownership when algorithmic outcomes cause loss or discrimination	Weak role attribution and incomplete decision provenance	Accountability Structure; Strategic Oversight Mechanism	Over-automation of managerial judgment	Managers defer reflexively to machine outputs even when a contextual override is warranted	Excessive delegation without calibrated human intervention	Algorithmic Decision Layer; Accountability Structure; Strategic Oversight Mechanism
Board-level symbolic oversight	AI appears on governance agendas, but directors lack interpretive capability and only ratify management	Oversight detached from technical substance	Strategic Oversight Mechanism; Adaptive Feedback and Audit Loop	Fragmented governance across silos	Legal, compliance, data science, and strategy units each govern AI separately without integration	No cross-layer coordination architecture	All SOA layers, especially Strategic Oversight Mechanism and Adaptive Feedback and Audit Loop

Reactive reputational management	Firms respond publicly to algorithmic controversy but fail to redesign internal controls	Governance is treated as external communication rather than internal architecture	Accountability Structures: Risk Management and Compliance Shields Adapted Feedback and Accountability Loop

More broadly, incorporating feedback loops fosters organizational reflexivity, enabling firms to move beyond reactive compliance toward anticipatory governance. By continuously aligning algorithmic systems with shifting regulatory expectations, societal norms, and strategic priorities, organizations position themselves to navigate the uncertainties of an algorithm-driven environment with resilience and legitimacy. In doing so, they transform governance from a constraint on innovation into a dynamic capability that supports sustainable value creation [12, 17].

Illuminating Scholarly Contributions: Advancing Digital Governance Theory Through SOAR

The SOAR Framework extends existing theory in three substantive ways, each addressing a documented lacuna in the literature.

Reframing accountability in hybrid systems

By formalizing role-based liability mapping across human-AI interfaces, SOAR moves beyond philosophical discussions of accountability [3, 11] toward an operational model that assigns clear ownership at every layer.

Elevating transparency as a strategic imperative

SOAR reframes transparency not as a regulatory burden but as a source of competitive differentiation and trust capital [4, 14, 15]. It integrates technical explainability with

organizational disclosure protocols, thereby bridging the gap between computer-science interpretability research and strategic management scholarship [18, 20].

Integrating multi-layer architectures into management studies

Prior frameworks addressed isolated elements—AI ethics guidelines, board competencies, or algorithmic management—yet none synthesized them into a dynamic, feedback-rich ecosystem [1, 5, 21]. SOAR fills this integration void, offering digital business scholars a unified lens for studying socio-technical governance in the algorithmic age [17, 23, 26]. Its boundary conditions (applicable primarily to medium-to-large organizations with mature data infrastructures) invite contextual extensions and comparative research.

Charting Future Trajectories: Envisioning Ethical and Strategic Horizons for Algorithmically Governed Enterprises

The SOAR Framework signals a paradigm shift from governance-as-control to governance-as-stewardship, where algorithmic systems become strategic partners rather than opaque black boxes.

Emerging research avenues for empirical validation

Future studies should test SOAR's efficacy through longitudinal case studies and quasi-experimental designs in sectors ranging from finance to healthcare [16, 19]. Key questions include: How do feedback-loop intensity levels moderate accountability outcomes? Does transparency reporting enhance or erode firm valuation? Such inquiries will refine the model and extend its applicability.

Policy recommendations for regulatory alignment

Regulators should incorporate SOAR-style multi-layer requirements into forthcoming AI acts, mandating algorithmic oversight committees and annual transparency disclosures [2, 7, 13]. This alignment would harmonize

corporate practice with public interest while preserving innovation incentives [21, 27-29].

Transformative potential for sustainable competitive advantage

Organizations that fully internalize SOAR will convert algorithmic mediation from a governance liability into a source of enduring advantage—ethical, resilient, and strategically agile. By institutionalizing accountability, enforcing transparency, and sustaining strategic oversight, they will lead the next wave of digitally responsible enterprise.

The SOAR Framework thus offers both a diagnostic mirror and a navigational compass for the algorithmically mediated organization. Its successful adoption will determine whether the promise of artificial intelligence translates into responsible, trustworthy, and strategically governed progress.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 15 Nov 2023 Revised: 05 Jan 2024 Accepted: 20 Feb 2024
Published online: 18 March 2024

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Mäntymäki M, Minkkinen M, Birkstedt T, Viljanen M. Defining organizational AI governance. *AI Ethics*. 2022;2(3):603-9.
- Hickman E, Petrin M. Trustworthy AI and Corporate Governance: The EU's Ethics Guidelines for Trustworthy Artificial Intelligence from a Company Law Perspective. *Eur Bus Org Law Rev*. 2021;22(4):593-625.
- Adner R. Ecosystem as structure: an actionable construct for strategy. *J Manage*. 2017;43(1):39-58.
- Jacobides MG, Cennamo C, Gawer A. Towards a theory of ecosystems. *Strateg Manage J*. 2018;39(8):2255-76.
- Birkstedt T, Minkkinen M, Tandon A, Mäntymäki M. AI governance: themes, knowledge gaps and future agendas. *Internet Res*. 2023;33(7):133-67.
- Buhmann A, Paßmann J, Fieseler C. Managing Algorithmic Accountability: Balancing Reputational Concerns, Engagement Strategies, and the Potential of Rational Discourse. *J Bus Ethics*. 2020;163(2):265-80.
- Hein A, Schreieck M, Riasanow T, Soto Setzke D, Wiesche M, Böhm M, et al. Digital platform ecosystems. *Electron Mark*. 2020;30(1):87-105.
- Mukhopadhyay S, Bouwman H. Orchestration and governance in digital platform ecosystems: a literature review and trends. *Digit Policy Regul Gov*. 2019;21(4):329-51.
- Schildt H. Big data and organizational design – the brave new world of algorithmic management and computer augmented transparency. *Innovation (N S)*. 2017;19(1):23-30.

Jarrahi MH, Newlands G, Lee MK, Wolf CT, Kinder E, Sutherland W. Algorithmic management in a work context. *Big Data Soc.* 2021;8(2):20539517211020332.

Lepri B, Oliver N, Letouzé E, Pentland A, Vinck P. Fair, Transparent, and Accountable Algorithmic Decision-making Processes. *Philos Technol.* 2018;31(4):611-27.

Benlian A, Wiener M, Cram WA, Krasnova H, Maedche A, Möhlmann M, et al. Algorithmic Management. *Bus Inf Syst Eng.* 2022;64(6):825-39.

Krafft TD, Zweig KA, König PD. How to regulate algorithmic decision-making: A framework of regulatory requirements for different applications. *Regul Gov.* 2022;16(1):119-36.

Janssen M, Brous P, Estevez E, Barbosa LS, Janowski T. Data governance: Organizing data for trustworthy Artificial Intelligence. *Gov Inf Q.* 2020;37(3):101493.

Larsson S, Heintz F. Transparency in artificial intelligence. *Internet Policy Rev.* 2020;9(2):1-16.

Haefner N, Wincent J, Parida V, Gassmann O. Artificial intelligence and innovation management: A review, framework, and research agenda. *Technol Forecast Soc Change.* 2021;162:120392.

Berente N, Gu B, Recker J, Santhanam R. Managing artificial intelligence. *MIS Q.* 2021;45(3):1433-50.

Li J, Li M, Wang X, Thatcher JB. Strategic directions for AI: The role of CIOs and boards of directors. *MIS Q.* 2021;45(3):1603-44.

Vomberg A, Schauerte N, Krakowski S, Ingram Bogusz C, Gijzenberg MJ, Bleier A. The cold-start problem in nascent AI strategy: Kickstarting data network effects. *J Bus Res.* 2023;168:114236.

Abdel-Karim BM, Pfeuffer N, Carl KV, Hinz O. How AI-based systems can induce reflections: The case of AI-augmented diagnostic work. *MIS Q.* 2023;47(4):1395-424.

Cihon P, Schuett J, Baum SD. Corporate Governance of Artificial Intelligence in the Public Interest. *Information.* 2021;12(7):275.

Ågerfalk PJ. Artificial intelligence as digital agency. *Eur J Inf Syst.* 2020;29(5):497-512.

Burström T, Parida V, Lahti T, Wincent J. AI-enabled business-model innovation and transformation in industrial ecosystems: A framework, model and outline for further research. *J Bus Res.* 2021;121:221-31.

Autio E. Orchestrating ecosystems: a multi-layered framework. *Innovation (N S).* 2022;24(1):96-109.

Marabelli M, Newell S, Handunge V. The lifecycle of algorithmic decision-making systems: Organizational choices and ethical challenges. *J Strateg Inf Syst.* 2021;30(3):101665.

Rai A, Constantinides P, Sarker S. Editor's comments: Next-generation digital platforms: Toward human-AI hybrids. *MIS Q.* 2019;43(1):iii-xii.

Linde L, Sjödin D, Parida V, Wincent J. Dynamic capabilities for ecosystem orchestration: A capability-based framework for smart city innovation initiatives. *Technol Forecast Soc Change.* 2021;166:120614.

de Reuver M, Sørensen C, Basole RC. The digital platform: a research agenda. *J Inf Technol.* 2018;33(2):124-35.

Addo A. Orchestrating a digital platform ecosystem to address societal challenges: A robust action perspective. *J Inf Technol.* 2022;37(4):359-86.