AGORA International Journal of Juridical Sciences, http://univagora.ro/jour/index.php/aijjs
ISSN 1843-570X, E-ISSN 2067-7677
Vol. 19, No. 2 (2025), pp. 408-418

ETHICS BY DESIGN: INTEGRATING LEGAL PRINCIPLES INTO THE MANAGEMENT OF ARTIFICIAL INTELLIGENCE

D. M. MINGHIRAȘI

Dragos Marius Minghirași
Università degli Studi della Campania „Luigi Vanvitelli", Italy
https://orcid.org/0009-0003-2639-3180, E-mail: minghirasidragosmarius@yahoo.com

Abstract: *Artificial intelligence (AI) has moved from experimental technology to an integral component of strategic management. Yet, as algorithms increasingly shape human decisions, the question of how to embed ethics and legality into AI systems becomes central. This article examines the concept of Ethics by Design — the proactive integration of ethical and legal principles into the architecture, governance, and operation of AI. By bridging law, management, and ethics, the study argues that sustainable innovation requires more than regulatory compliance: it demands organisational structures that internalise moral and legal accountability. The paper outlines the foundations of ethical AI governance, the managerial implications of algorithmic accountability, and the role of law in ensuring human-centred digital transformation.*

Keywords: *ethics by design, AI governance, legal responsibility, strategic management, accountability, sustainability*

1. Introduction

Artificial intelligence has transformed organisational management, enabling unprecedented efficiency and predictive capability. However, these benefits come with ethical and legal challenges: opacity, bias, data misuse, and diminished human oversight. In response, the principle of Ethics by Design has emerged as a guiding philosophy, advocating that ethical and legal standards must be integrated into the very design and deployment of AI systems — not added later as corrective mechanisms.

This article situates Ethics by Design at the intersection of law, management, and technology, proposing it as a managerial framework that aligns innovation with legitimacy. Rather than viewing law as a constraint, it considers legality and ethics as structural components of sustainable digital governance.

Building on this premise, Ethics by Design reframes artificial intelligence not as a neutral technical tool, but as a socio-technical system whose consequences reverberate across organisational, legal, and societal domains. In this sense, the concept moves beyond traditional compliance models. Instead of merely ensuring that AI systems do not violate existing rules, Ethics by Design encourages organisations to embed *anticipatory governance* into every stage of the AI lifecycle — from data collection and model training to deployment, monitoring, and decommissioning.

This proactive orientation is crucial in a landscape where regulatory frameworks are evolving rapidly. Instruments such as the EU Artificial Intelligence Act, the General Data Protection Regulation (GDPR), and sector-specific compliance requirements impose not only

technical obligations but also managerial responsibilities. Organisations are increasingly required to demonstrate that they exercise *due diligence*, maintain *risk-based governance structures*, and uphold *transparency obligations*. Ethics by Design provides the conceptual and operational tools necessary to meet these standards, transforming abstract legal principles into concrete managerial practices.

Within organisational management, the integration of Ethics by Design operates on three interconnected levels:

1. Strategic level – redefining corporate governance to incorporate ethical risk-assessment, accountability structures, and decision-making principles aligned with fundamental rights and societal values.
2. Operational level – translating ethical frameworks into procedures, workflows, quality controls, and documentation practices that ensure traceability and accountability throughout the AI lifecycle.
3. Cultural level – fostering an organisational mindset where employees, managers, and technical teams internalise ethical reflexivity as part of everyday practice, not as an external constraint.

Through this multilevel approach, Ethics by Design becomes a managerial philosophy capable of bridging the gap between *legal compliance*, *technological innovation*, and *institutional legitimacy*. In a world where algorithmic systems increasingly mediate hiring decisions, credit scoring, law enforcement tools, medical diagnostics, and educational evaluation, the question is no longer whether organisations should adopt ethical frameworks, but how effectively they can integrate them into their governance structures.

Ultimately, the argument advanced in this article is that Ethics by Design is not merely a normative ideal — it is an operational necessity. Organisations that treat ethics as a peripheral or optional consideration expose themselves to reputational, legal, and operational risks. Conversely, those that embed ethical principles into the architecture of their AI systems cultivate trust, legitimacy, and long-term sustainability.

2. Conceptual Framework

2.1. From Compliance to Ethical Design

Traditional compliance models address legal risks after technology has been implemented. Ethics by Design inverts this logic: it seeks to embed moral and legal reasoning from the earliest stages of AI development. This paradigm shift moves organisations from reactive governance to proactive accountability.

The European Union's AI Act (2024) embodies this transition by requiring that systems meet principles of transparency, fairness, and human oversight before deployment. Management, therefore, becomes a co-creator of ethical architecture rather than a passive regulator.

This shift reflects a broader evolution in governance theory: the recognition that compliance alone is insufficient in environments characterised by technological opacity and rapid innovation cycles. Traditional models assume that risks can be identified, measured, and sanctioned ex post, once a violation occurs. But algorithmic systems challenge this assumption. Their complexity, reliance on probabilistic reasoning, and dependence on dynamic datasets create forms of uncertainty that cannot be adequately addressed through post-hoc remedies.

Ethics by Design introduces a preventive logic, seeking to eliminate or mitigate risks *before* they materialise. In doing so, it repositions ethical and legal assessments as integral components of the design process. Engineers, data scientists, and organisational leaders are encouraged to collaborate from the outset, ensuring that system architecture reflects normative commitments such as non-discrimination, proportionality, explainability, and respect for fundamental rights. Under this framework, management assumes a transformative role. Rather than merely enforcing compliance checklists, managers act as custodians of organisational values and stewards of responsible innovation. Their involvement is essential for aligning technical choices with broader institutional goals — including reputation, stakeholder trust, and long-term sustainability.

The EU AI Act formalises this managerial responsibility by mandating risk-management systems, documentation practices, human-in-the-loop controls, and continuous monitoring *prior* to market entry. The law effectively codifies the Ethics by Design logic, signalling that ethical integration is no longer optional but a structural requirement for digital governance in Europe. This represents a profound change: the burden is now on organisations to demonstrate that their systems were developed within a principled, rights-respecting framework.

In this sense, Ethics by Design functions not merely as a compliance strategy, but as a holistic governance model. It cultivates anticipatory accountability, reshaping organisational processes so that ethical reflection becomes embedded in the everyday functioning of AI development and deployment. This approach enhances institutional legitimacy and prepares organisations for emerging regulatory landscapes, where societal expectations increasingly demand transparency, fairness, and human-centric oversight.

2.2. The Legal Foundations of AI Governance

AI governance rests on existing legal doctrines—data protection, human rights, and liability law—but extends them into uncharted terrain. Concepts such as algorithmic accountability and explainability represent managerial translations of legal principles like due process and non-discrimination. In this sense, legal governance and managerial strategy converge, ensuring that ethical compliance is not a legal checkbox but a performance criterion.

AI governance is therefore anchored in a multilayered legal ecosystem, where classical legal doctrines intersect with emerging regulatory instruments. While data protection law (notably the GDPR) establishes foundational requirements for lawful processing, transparency, and individual rights, human rights law frames the broader normative boundaries within which AI systems must operate. Liability law, in turn, delineates the conditions under which organisations may be held accountable for harm caused by algorithmic decision-making.

Yet these frameworks were not originally designed with autonomous, adaptive, or opaque technologies in mind. As a result, AI governance evolves by extending established legal concepts into new domains, requiring both reinterpretation and innovation. Algorithmic accountability—requiring organisations to justify and document the internal logic of AI systems—translates the procedural fairness of due process into a technological context. Similarly, explainability operationalises the right to understanding and contestation, bridging

the gap between complex computational reasoning and the individual's capacity to challenge decisions that affect them.

These developments illustrate a growing convergence between legal governance and managerial strategy. Compliance is no longer confined to legal departments; it becomes a shared organisational responsibility embedded in design decisions, procurement policies, risk management frameworks, and oversight mechanisms. In this context, legal norms act as strategic constraints and enablers: they shape how organisations allocate resources, design workflows, and evaluate performance.

Thus, ethical compliance is reframed as a performance criterion—a measurable indicator of organisational maturity and resilience. It impacts stakeholder trust, regulatory approval, market competitiveness, and long-term sustainability. Far from being a bureaucratic checkbox, legal adherence becomes integral to innovation itself: an organisation that cannot demonstrate accountability, transparency, or fairness risks exclusion from markets, erosion of public trust, and exposure to financial and reputational sanctions.

In this way, the legal foundations of AI governance provide the normative architecture upon which managerial frameworks such as Ethics by Design are built. They ensure that technological ambition is aligned with societal values and that organisations internalise their obligations not as defensive reactions to regulation, but as proactive commitments to legitimacy and responsible digital leadership.

2.3. Governance Through Design: Embedding Law into Technical Architecture

The evolution from reactive compliance to proactive ethical integration requires a fundamental reconceptualisation of governance mechanisms. Instead of relying on external audits, periodic checks, or ex post sanctioning, Governance Through Design embeds legal and ethical principles directly into the technical architecture of AI systems. This approach treats algorithms not merely as computational tools, but as infrastructures of normative significance — systems whose very structure can promote or undermine fundamental rights.

Embedding governance into design means translating abstract legal norms into operational specifications. Requirements such as data minimisation, purpose limitation, fairness, explainability, and human oversight must be encoded into system architectures, not merely referenced in policy documents. This technical translation process elevates the role of multidisciplinary collaboration: engineers must understand legal obligations, lawyers must grasp technical constraints, and managers must coordinate these perspectives into a coherent organisational strategy.

Within this paradigm, risk management becomes a design principle rather than an administrative afterthought. High-risk AI systems, as defined by the EU AI Act, must incorporate mechanisms for continuous monitoring, auditability, traceability of decisions, and the capacity to intervene—features that transform governance from an external control into an internal function of the system itself. Technical features such as logging, version control, explainability layers, and human-in-the-loop decision gates are no longer optional enhancements but mandatory safeguards that embody the organisation's commitment to lawful and ethical operation. This approach also reflects a broader theoretical movement in regulatory thinking: the shift from *rules-based* to *architecture-based* governance. Whereas traditional regulation relies on setting behavioural expectations, governance through design builds

constraints and protections into the environment in which decisions are made. In the context of AI, this means ensuring that systems are structured to prevent discriminatory outcomes, avoid opacity, and preserve human agency by default.

For managers, this transformation blurs the line between strategy and compliance. Ethical governance becomes a core organisational competence, requiring investment in training, infrastructure, and interdisciplinary teams. It also demands cultural change, as employees must internalise the idea that every design decision—however technical—carries normative implications. In this sense, Governance Through Design represents the operational heart of Ethics by Design: it turns values into features, principles into protocols, and legal obligations into system behaviour.

By embedding governance into the very code and architecture of AI systems, organisations simultaneously enhance legal robustness, ethical credibility, and operational resilience. The result is a governance model capable not only of meeting regulatory expectations, but of shaping a sustainable and rights-centric technological ecosystem.

3. Management of Artificial Intelligence: Ethical and Organisational Dimensions

3.1. The Manager as Ethical Designer

The role of the manager evolves from operational control to ethical stewardship. Effective AI management requires multidisciplinary teams that combine technical expertise with legal insight and ethical reasoning. Decision-makers must implement internal review boards, algorithmic audit systems, and risk assessment tools to ensure that AI aligns with both regulatory norms and social expectations. Embedding ethical thinking into corporate governance structures — through codes of conduct, training, and board-level oversight — transforms Ethics by Design into a continuous management process rather than a one-time design feature.

In this emerging governance landscape, the manager is no longer a passive overseer of technical processes but an ethical architect responsible for orchestrating the conditions under which AI systems can operate responsibly. This redefined role requires an ability to navigate the intersection of law, technology, and organisational behaviour — a hybrid competence that blends strategic foresight with normative sensitivity. Managers must cultivate a working environment in which ethical reflection is not peripheral but embedded in the organisational DNA.

To achieve this, multidisciplinary collaboration becomes a structural requirement rather than an optional enhancement. Ethical AI design depends on the interaction of data scientists, legal experts, ethicists, cybersecurity specialists, human resources professionals, and executive leadership. The manager's task is to facilitate this collaboration, align diverse forms of expertise, and ensure that ethical deliberation informs every stage of the AI lifecycle. This includes setting priorities, allocating resources, defining success metrics, and ensuring transparency in decision-making processes.

A central managerial responsibility is the establishment of internal governance mechanisms that institutionalise ethical oversight. These include internal review boards for high-risk AI systems, algorithmic audit protocols, impact assessments, and mechanisms for reporting or mitigating harm. Such structures function as early-warning systems, enabling

organisations to detect biases, identify systemic risks, and adapt to regulatory changes before issues escalate into legal or reputational crises.

Furthermore, managers must embed ethical thinking into corporate governance frameworks. This can be achieved through codes of conduct tailored to AI development, continuous training programs on legal and ethical issues, and the integration of AI governance into board-level agendas. Board oversight—particularly through risk committees or dedicated ethics committees—ensures that organisational strategy aligns with societal values and regulatory expectations. In this context, Ethics by Design is not a static requirement but a continuous management process, informed by iterative learning, monitoring, and adaptation.

This ongoing commitment to ethical stewardship enhances organisational legitimacy. It demonstrates to regulators, customers, employees, and partners that the organisation takes seriously its responsibility to manage AI systems in ways that respect human rights and promote societal well-being. Ultimately, the manager as ethical designer becomes a pivotal actor in shaping a form of digital governance that is not only technically robust but also morally grounded.

3.2. Accountability in Algorithmic Decision-Making

Algorithmic accountability challenges the classical notion of responsibility in both law and management. Who is responsible when an algorithm discriminates or causes harm? Ethics by Design responds by distributing accountability across the entire lifecycle of AI: design, deployment, and monitoring. This distributed model of responsibility reflects principles of good governance — transparency, traceability, and due diligence — adapted to the digital context.

Algorithmic accountability thus requires a reconceptualisation of responsibility that departs from traditional, human-centred models. Classical legal doctrine presupposes identifiable actors who intentionally or negligently cause harm. AI systems, however, complicate this framework: they operate autonomously, evolve over time, and rely on datasets and optimisation processes that often obscure causal chains. This raises critical questions for both law and management: How do we attribute fault in environments where decision-making is partially automated, probabilistic, and opaque?

Ethics by Design offers a structural response by distributing accountability across the entire AI lifecycle. Rather than assigning responsibility exclusively to developers or end-users, this model recognises that multiple actors contribute to algorithmic outcomes. Designers influence model architecture, managers shape governance structures, procurement teams select datasets and vendors, and operators oversee deployment and maintenance. Accountability therefore becomes multi-layered, shared, and process-based, aligning with contemporary understandings of organisational responsibility.

This distributed approach operationalises core principles of good governance:

- Transparency ensures that decision-making processes are documented, explainable, and open to scrutiny by regulators, stakeholders, and affected individuals.
- Traceability provides the capacity to reconstruct the decision-making pathway, enabling organisations to identify where biases, errors, or system failures emerged.

- Due diligence demands continuous risk assessments, monitoring, and corrective actions, signalling that responsibility is not static but evolves as systems learn and adapt.

Importantly, this reconfiguration of accountability does not dilute responsibility; rather, it enhances organisational robustness by ensuring that no stage of AI development is exempt from oversight. It shifts accountability from an ex post attribution of blame to an ex ante commitment to responsible design and governance, encouraging organisations to anticipate and mitigate potential harms before they materialise.

From a managerial perspective, this means establishing clear lines of responsibility within the organisation: documenting roles and expectations, creating escalation pathways for ethical concerns, and integrating algorithmic risk into performance evaluation and corporate reporting. This systemic accountability not only meets emerging regulatory standards — such as those embedded in the EU AI Act — but also strengthens institutional legitimacy by demonstrating a proactive commitment to rights-respecting digital innovation.

Ultimately, algorithmic accountability becomes a cornerstone of Ethics by Design: it reframes responsibility not as a punitive mechanism, but as a collective, ongoing practice of stewardship, ensuring that AI systems function in ways that align with societal values and legal obligations.

4. Organisational Governance and Legal Integration

4.1. Embedding Legal Principles in Corporate Strategy

To operationalise Ethics by Design, organisations must integrate legal principles into strategic planning. Legal counsel should participate in AI strategy meetings, and ethical KPIs (Key Performance Indicators) should be embedded in management scorecards. This convergence of legal and managerial oversight enhances institutional resilience, helping organisations anticipate legal risks while maintaining reputational integrity.

Embedding legal principles into corporate strategy requires organisations to move beyond the episodic consultation of legal experts and toward a model where legal reasoning becomes a strategic asset. In the context of AI governance, this means that legal counsel, compliance officers, and data protection specialists must be included from the inception of technological initiatives, not merely engaged at the final approval stage. Early integration allows legal principles — such as proportionality, fairness, accountability, and respect for fundamental rights — to shape system architecture, risk appetite, and organisational priorities.

In practical terms, organisations can operationalise this alignment by incorporating ethical and legal performance indicators into managerial scorecards. These KPIs may include metrics related to algorithmic bias detection, explainability audits, data quality assurance, incident reporting timelines, or the rate of compliance with internal governance protocols. When ethical performance becomes measurable, it also becomes manageable, enabling leaders to evaluate the organisation's progress and identify areas requiring further intervention.

Moreover, this convergence of legal and managerial oversight contributes to institutional resilience. By embedding legal considerations into strategic planning, organisations develop the capacity to anticipate regulatory shifts, adapt to evolving societal expectations, and respond to emerging risks before they materialise. This anticipatory posture

reduces exposure to litigation, regulatory penalties, and market disruption — consequences that can be particularly severe in high-risk AI domains such as healthcare, finance, education, or law enforcement. At the same time, integrating legal principles at the strategic level strengthens reputational integrity. In an era where stakeholder trust is increasingly tied to responsible data practices and transparent algorithmic governance, organisations that can demonstrate principled decision-making enjoy competitive advantages. Ethical credibility becomes part of the brand identity, reinforcing consumer confidence, investor trust, and employee engagement.

Thus, embedding legal principles into corporate strategy is not merely a compliance exercise but a strategic imperative. It ensures that Ethics by Design is translated into actionable governance structures, long-term institutional vision, and operational coherence. By aligning legal norms with managerial objectives, organisations create a governance ecosystem in which innovation and legitimacy reinforce each other, rather than standing in tension.

4.2. The Role of Law as an Enabler, not a Constraint

While many firms view regulation as an impediment to innovation, this article argues that legal frameworks — particularly those centred on human rights and data protection — provide a foundation for sustainable technological progress. Incorporating ethics by design ensures not only compliance but trustworthiness, which in turn strengthens market credibility and public acceptance.

Contrary to the perception that regulation stifles innovation, law functions as an enabling architecture that provides clarity, stability, and legitimacy for technological development. In domains characterised by rapid innovation and high uncertainty, such as artificial intelligence, legal frameworks reduce ambiguity by articulating the boundaries within which innovation can occur responsibly. Far from constraining creativity, these boundaries establish the *conditions of possibility* for sustainable technological ecosystems.

Human rights law, data protection norms, and emerging AI-specific regulations offer normative guardrails that guide organisations toward solutions that are socially acceptable and ethically sound. They encourage developers to consider impacts on autonomy, dignity, non-discrimination, and privacy — principles that are essential for maintaining public trust in automated decision-making. Legal requirements such as transparency, fairness, and human oversight incentivise firms to adopt methodological rigour and operational discipline, thereby improving the quality and robustness of AI systems.

Understanding law as an enabler also reframes compliance as a strategic advantage rather than a bureaucratic obligation. Organisations that embed legal principles into their AI governance models are better positioned to avoid costly redesigns, regulatory sanctions, or reputational crises. They can enter markets with greater confidence, secure partnerships more easily, and differentiate themselves through ethical leadership. In high-stakes industries, where consumer and stakeholder scrutiny is intense, demonstrable adherence to legal and ethical standards becomes a form of *competitive capital*.

Ethics by Design amplifies this enabling function by ensuring that normative considerations are integrated from the outset. This proactive incorporation fosters trustworthiness, a quality increasingly recognised as essential for the societal acceptance of AI technologies. Trust is not merely an intangible value; it has concrete market implications —

influencing user adoption rates, customer loyalty, investor confidence, and regulatory goodwill. Companies that can demonstrate the integrity of their systems enjoy a reputational premium, especially in sectors where the consequences of algorithmic harm are significant.

Thus, law emerges not as an obstacle but as a catalyst for sustainable innovation. By embedding legal principles into design processes and organisational strategy, firms create AI systems that are not only technologically advanced but also aligned with societal expectations and democratic values. This alignment strengthens the legitimacy of AI deployment and supports a vision of technological progress that is both responsible and resilient.

5. The Future of Ethical AI Governance

5.1. Human-Centred Design and the "Right to Explanation"

A core challenge in AI governance is maintaining human agency. The EU General Data Protection Regulation (GDPR) introduced the "right to explanation", obliging organisations to clarify automated decisions that affect individuals. Managers must therefore ensure that AI systems remain interpretable and controllable — maintaining the principle that technology serves humanity, not vice versa.

The rise of automated decision-making has intensified concerns about the erosion of human agency. As algorithms increasingly shape access to employment, credit, healthcare, education, and public services, the need to preserve individual autonomy becomes a central pillar of AI governance. Human-centred design responds to this challenge by ensuring that technological systems enhance, rather than displace, human judgment, dignity, and self-determination.

Within this framework, the "right to explanation", emerging from the GDPR and echoed in subsequent European regulatory instruments, represents a crucial safeguard. Although debated in its precise legal scope, the principle affirms that individuals affected by automated decisions are entitled to understand the reasoning behind them and to contest outcomes that are unfair or erroneous. This requirement reflects deeper legal values — procedural fairness, transparency, and respect for individual rights — now translated into the digital environment.

For organisations, the right to explanation has profound managerial implications. It demands that AI systems be designed with interpretability, traceability, and controllability at their core. Managers must ensure that technical teams adopt methods that allow for meaningful insights into how decisions are generated, whether through inherently interpretable models, explainability layers, or rigorous documentation practices. This design orientation counters the "black box" problem, which undermines accountability and erodes public trust.

Moreover, maintaining human agency requires effective human oversight mechanisms. Managers must define clear thresholds for human intervention, establish escalation routines when automated decisions produce anomalies, and ensure that employees are trained to understand the capabilities and limitations of AI tools. Such safeguards are integral to preserving the principle that technology serves humanity, not vice versa.

Human-centred design therefore becomes a strategic priority, not a peripheral ethical concern. It aligns organisational practices with both regulatory mandates and societal expectations, reinforcing trust and legitimacy in AI-driven processes. By embedding

interpretability and oversight into technological infrastructures, managers uphold the core democratic principle that individuals retain the ability to understand, challenge, and influence the systems that shape their lives.

Ultimately, the right to explanation is not merely a procedural entitlement; it is a manifestation of a broader commitment to human-centred governance. As organisations increasingly rely on algorithmic tools, this commitment ensures that innovation remains aligned with human values, ethical constraints, and legal protections — the essential foundations of a responsible digital society.

5.2. The Sustainability Dimension

Ethical AI governance contributes directly to organisational sustainability. Trust, transparency, and fairness enhance long-term stakeholder value, reduce litigation risk, and foster innovation ecosystems built on legitimacy. In this light, Ethics by Design is not a moral accessory but a strategic asset — an element of competitive differentiation in the digital economy.

Sustainability in the digital era extends far beyond environmental concerns. It encompasses the capacity of organisations to operate responsibly, maintain stakeholder trust, and remain resilient amid technological, legal, and social change. In this broader sense, ethical AI governance is inseparable from organisational sustainability. AI systems that violate privacy, produce discriminatory outcomes, or operate opaquely undermine not only regulatory compliance but also institutional credibility, social licence to operate, and long-term value creation.

Ethics by Design strengthens this sustainability framework by embedding principles such as fairness, transparency, and accountability into organisational processes. These principles function as risk-mitigating mechanisms, reducing exposure to litigation, regulatory sanctions, reputational harm, and operational disruptions. As AI-driven decisions increasingly affect fundamental human interests, stakeholders — including consumers, employees, regulators, investors, and civil society — expect organisations to demonstrate ethical integrity. Meeting these expectations is essential for securing durable relationships and maintaining a stable operating environment.

In addition, ethical AI governance promotes innovation ecosystems grounded in legitimacy. When organisations adopt clear ethical standards, they create predictable and trustworthy environments for technological experimentation and cross-sector collaboration. Partners and developers are more willing to engage with institutions that demonstrate transparency and responsible governance, while customers show higher adoption rates for technologies perceived as safe, explainable, and aligned with societal values. This reinforces the idea that trust is not a soft value but a measurable component of economic performance.

From a strategic standpoint, Ethics by Design becomes a source of competitive differentiation in the digital economy. As regulatory standards tighten and societal scrutiny intensifies, organisations that proactively integrate ethical and legal considerations gain a market advantage. They can enter regulated markets more quickly, secure certification more easily, and mitigate the costly need for retrofitted compliance measures. In contrast, firms that treat ethics as an afterthought risk technological lock-in, systemic vulnerabilities, and loss of stakeholder confidence.

Thus, the sustainability dimension reveals Ethics by Design as far more than a moral accessory. It is a strategic asset—a driver of resilience, innovation, and long-term institutional legitimacy. By embedding ethical governance into organisational culture and technical infrastructure, managers cultivate systems capable not only of technical performance but of enduring societal acceptance. In a digital landscape shaped by rapid disruption and heightened expectations, this alignment between ethics and sustainability becomes indispensable for responsible organisational leadership.

6. Conclusions and Managerial Implications

The management of artificial intelligence requires a synthesis of law, ethics, and strategy. Ethics by Design transforms this synthesis into a practical governance model, ensuring that technological innovation aligns with human values and legal standards. Organisations that institutionalise ethical design gain resilience, legitimacy, and sustainability — not because they follow rules, but because they internalise them into culture and process.

The future of management in the AI era will belong to leaders who understand that ethics is not external to design but the very framework through which intelligent systems become truly responsible.

REFERENCES

1. European Commission (2024). Artificial Intelligence Act. Brussels: EU Publications.
2. Floridi, L. (2020). The Logic of Information: A Theory of Philosophy as Conceptual Design. Oxford University Press.
3. Mittelstadt, B. (2022). “Principles Alone Cannot Guarantee Ethical AI.” Nature Machine Intelligence, 4(3), 187–190.
4. OECD (2023). AI Principles and Governance Framework. Paris: OECD Publishing.
5. Tranter, K. (2021). Living in Technical Legality: Science Fiction and Law as Technology. Edinburgh University Press.
6. World Economic Forum (2022). Responsible AI: A Framework for Ethical Management. Geneva: WEF Publications.