



HUMANISTIC LEADERSHIP IN THE AGE OF ARTIFICIAL INTELLIGENCE: A HUMAN-CENTRED GOVERNANCE AND MEASUREMENT FRAMEWORK FOR AI-ENABLED ORGANIZATIONS

Ignacio Bonasa Alzuria
President and Founder,
University of Happiness and Well-Being;
President and Co-Founder,
Liderarte Factoría de Talento, Spain
ORCID: 0009-0001-3940-4278
ibonasa@hotmail.com

Corresponding author: Ignacio Bonasa Alzuria

DOI Link: <https://doi.org/10.64971/j.cph.eijtem.v13.i3.9.2026>

ABSTRACT

Artificial intelligence (AI) is expanding organizational analytical capacity while redistributing decision power across models, data, interfaces and vendors. This article examines how leadership can govern that redistribution without reducing responsible AI to technical control or regulatory compliance. Using a structured integrative review and conceptual theory-building approach, it synthesizes humanistic management, human-centred AI, responsible AI governance, meaningful work, psychological safety and Industry 5.0. The article develops the Humanistic Leadership in Artificial Intelligence (HLAI) Framework, an eight-domain governance architecture linking technological capability to human, organizational and societal outcomes. Its original contribution lies in three integrations that existing frameworks generally treat separately: dignity and meaningful work as design criteria; leadership behaviour and decision rights as implementation mechanisms; and a proposed Humanistic AI Leadership Index (HAILI) as a measurement architecture. The framework is positioned explicitly against NIST AI RMF, ISO/IEC 42001, Shneiderman's human-centred AI and Industry 5.0, showing that HLAI complements rather than duplicates them. Twelve empirically testable propositions and a staged implementation roadmap are offered. The study does not claim psychometric validation; instead, it establishes a transparent basis for subsequent scale development and comparative research in engineering and technology management.

Keywords: humanistic leadership; artificial intelligence; HLAI Framework; Humanistic AI Leadership Index; responsible AI; augmented judgment; Industry 5.0; organizational flourishing; ethical governance; human-AI collaboration

1. INTRODUCTION

Artificial intelligence has moved from the margins of organizational life to its operating core. Algorithms recommend candidates, forecast demand, classify risks, generate content, monitor processes, support clinical and financial judgments, personalize learning and shape the information through which managers interpret reality. Generative AI has accelerated this transition because it places sophisticated computational capability directly in the hands of ordinary employees. The result is not merely another wave of automation. It is a reconfiguration of how knowledge is produced, authority is exercised and responsibility is distributed.

This technological acceleration arrives at a moment of institutional fragility. Many employees already experience work as overloaded, fragmented and emotionally thin. Trust in organizations is easily damaged by opaque decisions, constant surveillance, inconsistent values and the feeling that people are treated as interchangeable units of productivity. In such a context, introducing AI without a renewed philosophy of leadership can amplify existing weaknesses. A biased process can scale faster; a careless decision can travel further; and an impersonal culture can become embedded in interfaces, dashboards and automated workflows.

The central challenge is therefore not whether organizations should adopt AI, but what kind of organization they wish to become through its adoption. Technology is never introduced into a moral vacuum. Every system reflects choices about what counts, whose interests matter, which errors are tolerated and who has the right to question

a result. Leadership gives those choices direction. When leadership is dominated by short-term efficiency, AI is likely to be used primarily to reduce cost, intensify control and transfer risk. When leadership is humanistic, AI can be used to remove drudgery, expand access, strengthen judgment and create more meaningful forms of contribution.

Humanistic leadership begins from a simple but demanding premise: every person possesses intrinsic worth that is not exhausted by performance, compliance or market value. This premise does not oppose ambition, technology or economic results. It changes their order. Performance becomes sustainable when it serves a legitimate purpose; technology becomes valuable when it enlarges human capability; and profit becomes responsible when it is created without sacrificing dignity, fairness or the future. Humanistic leadership therefore offers a relevant response to the deepest governance questions of the AI era.

This article develops a conceptual and practical framework for that response. It integrates insights from humanistic management, responsible AI, stakeholder theory, psychological safety, meaningful work and organizational flourishing. Rather than presenting AI as either salvation or threat, it treats it as an amplifier. AI magnifies the quality of the goals, data, incentives and cultures into which it is introduced. The decisive variable is not technological sophistication alone, but the moral and relational maturity of the organization.

The article makes five contributions. First, it clarifies why AI intensifies the need for humanistic leadership. Second, it introduces the HLAI Framework and operationalizes it through the HLAI governance architecture, an eight-dimensional architecture for values, decision rights and organizational practices. Third, it differentiates humanistic AI leadership from traditional, digital and efficiency-centred AI leadership. Fourth, it proposes the Humanistic AI Leadership Index (HALI), a multidimensional measurement architecture that can be validated empirically. Fifth, it translates the framework into governance routines, engineering and technology-management applications, a phased roadmap and testable propositions.

2. RESEARCH PROBLEM, QUESTIONS AND OBJECTIVES

The research problem arises from a widening gap between the speed of AI deployment and the maturity of the leadership systems governing it. Organizations can acquire powerful tools faster than they can build shared norms for their responsible use. Technical teams may optimize accuracy while overlooking meaning; senior leaders may approve systems they cannot adequately question; and employees may be expected to collaborate with AI without clarity about accountability, learning or job redesign. This gap creates strategic, ethical and human risks.

The guiding research question is: How can humanistic leadership enable organizations to use artificial intelligence in ways that improve performance while protecting dignity, agency, trust and shared flourishing?

RQ1. Which characteristics of AI most strongly transform the responsibilities of organizational leaders?

RQ2. Which principles distinguish humanistic leadership from merely compliant or efficiency-centred AI management?

RQ3. How can these principles be converted into an actionable governance and operating model?

RQ4. Which outcomes and indicators can be used to evaluate whether human-AI collaboration is genuinely human-centred?

Accordingly, the objectives are to synthesize relevant theoretical perspectives; identify leadership tensions generated by AI; formulate an original integrative framework; propose organizational practices and metrics; and define a future research agenda suitable for technology, engineering and management contexts.

3. METHODOLOGY: INTEGRATIVE CONCEPTUAL REVIEW AND THEORY BUILDING

This article uses an integrative conceptual review rather than an empirical design. That choice is deliberate. The purpose is not to claim causal effects from a newly collected dataset, but to connect bodies of knowledge that are often treated separately and to derive a coherent leadership architecture. Integrative reviews are particularly useful when an emerging phenomenon crosses disciplinary boundaries and requires conceptual synthesis, critical interpretation and model development.

The analysis was organized around five literature domains: humanistic management and dignity; leadership and moral agency; responsible and trustworthy AI; psychological safety and meaningful work; and stakeholder-oriented organizational well-being. Seminal contributions were combined with recent policy and scholarly work on generative AI. Sources were evaluated for conceptual relevance, authority and applicability to organizational decision-making. The review did not treat promotional claims by technology vendors as independent evidence.

The theory-building process followed four stages. First, the article identified the distinctive organizational properties of AI, including opacity, scale, probabilistic output, data dependence, automation bias and rapid diffusion. Second, it mapped the leadership tensions produced by those properties. Third, it synthesized humanistic principles capable of addressing the tensions. Fourth, it translated those principles into dimensions, practices, questions and indicators. The resulting framework should therefore be understood as a testable conceptual model and a managerial design tool, not as a validated psychometric instrument.

Several safeguards were adopted to preserve intellectual integrity. No empirical sample, statistical result or field intervention is fabricated. Claims are bounded by the evidence and by the conceptual nature of the study. The framework is presented as an original synthesis whose propositions require future empirical testing across industries, cultures and organizational sizes.

3.1 Review scope and transparent selection logic

The review was conducted as a structured integrative synthesis rather than as a systematic review claiming exhaustive coverage. Sources were selected when they met at least one of four criteria: conceptual relevance to human dignity and leadership; direct relevance to AI governance or human-AI interaction; institutional authority in standards or regulation; or empirical relevance to work design, psychological safety and occupational transformation. Priority was given to peer-reviewed scholarship and primary institutional sources. Foundational works were retained where they supplied constructs that remain central, while recent regulatory and governance documents were incorporated to reflect the rapidly changing AI environment.

3.2 Analytical coding and framework construction

The selected literature was coded around recurring constructs: dignity, autonomy, purpose, accountability, transparency, contestability, fairness, psychological safety, augmentation, capability development, meaningful work and distributed value. These constructs were compared across disciplinary traditions and converted into design requirements. The eight dimensions of the HLAI Framework were retained only when they satisfied three tests: conceptual distinctiveness, practical observability and relevance across more than one sector. This process produced a theory-to-practice chain from philosophical premise, to leadership behaviour, to governance mechanism, to measurable outcome.

3.3 Quality boundaries

The method does not establish prevalence, effect sizes or causal superiority. It offers analytical generalization: a coherent set of propositions and operational categories that can be examined through future qualitative, experimental, longitudinal and multilevel research. This boundary is important because conceptual novelty should not be confused with empirical validation. The article therefore distinguishes clearly between established evidence, reasoned synthesis and propositions requiring testing.

4. WHY AI CHANGES THE WORK OF LEADERSHIP

4.1 From information scarcity to judgment scarcity

Traditional management systems were built in conditions where obtaining, processing and communicating information consumed substantial time. AI reduces many of those constraints. It can summarize thousands of documents, detect patterns invisible to unaided analysis and generate plausible options in seconds. Yet abundance does not eliminate the need for judgment. It changes its location. Leaders must decide which questions deserve attention, what evidence is sufficient, which trade-offs are acceptable and when a technically strong answer is socially or morally wrong.

The scarce resource becomes wise judgment: the capacity to integrate facts, context, values, experience and consequences. An AI system can estimate probability; it cannot bear responsibility. It can generate reasons; it does not possess a life story, vulnerability or moral stake in the outcome. Leadership in the AI era therefore requires stronger epistemic humility. Leaders must distinguish prediction from truth, fluency from understanding and optimization from legitimacy.

4.2 From visible hierarchy to distributed algorithmic power

Power in organizations has traditionally been associated with formal positions. AI distributes power through models, datasets, thresholds, prompts, interfaces and default settings. A data engineer may shape promotion opportunities; a vendor may influence credit decisions; and a product manager may determine which human behaviours are rewarded or made invisible. Because this power is dispersed, governance cannot remain confined to the boardroom or technology department.

Humanistic leadership makes algorithmic power visible. It clarifies who designs, approves, monitors, challenges and retires systems. It also recognizes that those affected by a system possess knowledge that designers may lack. Participation is not a ceremonial consultation after deployment. It is a source of risk intelligence and moral legitimacy.

4.3 From bounded errors to errors at scale

Human decisions can be inconsistent, biased and slow. AI can reduce some of these weaknesses, but it can also reproduce them with extraordinary consistency and reach. A poor managerial judgment may harm one team; a flawed model can affect thousands of people before anyone understands the pattern. Scale converts small design assumptions into institutional consequences.

This feature changes the meaning of leadership accountability. Responsible leaders do not ask only whether a system works on average. They ask who bears the cost when it fails, whether the error can be appealed, whether

the harm is reversible and whether vulnerable groups are disproportionately exposed. The moral seriousness of AI governance increases with its scale, speed and irreversibility.

4.4 From tasks to identities

AI adoption is often described as a portfolio of tasks to automate or augment. Employees, however, experience work not only as tasks but as identity, status, community and meaning. Removing repetitive work may be liberating, but removing the activities through which a professional feel competent and useful can be disorienting. Job redesign therefore requires more than process engineering.

Humanistic leaders attend to the emotional and identity transitions created by AI. They explain what will change, what will remain distinctly human, how people will learn and how contribution will be recognized. They do not demand enthusiasm as a condition of belonging. They create legitimate spaces for concern, experimentation and disagreement.

4.5 From static competence to continuous co-learning

The useful life of skills is shortening as tools evolve. Organizations cannot respond through occasional training alone. They need a culture of continuous co-learning in which employees learn to use AI, critique it, teach one another and redesign work. This learning includes technical literacy, but also ethical reasoning, domain judgment and the courage to stop an unsafe process.

Leaders model this posture when they disclose uncertainty, test their own assumptions and treat errors as opportunities for system improvement rather than occasions for blame. Psychological safety becomes an infrastructure for technological reliability. Without it, employees conceal failures, comply with doubtful outputs or avoid experimentation.

5. HUMANISTIC LEADERSHIP: PHILOSOPHICAL AND ORGANIZATIONAL FOUNDATIONS

5.1 Dignity as the first design principle

Humanistic management rejects the reduction of persons to instruments. In organizational practice, dignity means that people are treated as subjects capable of voice, choice and development. This has direct implications for AI. Systems affecting employment, health, credit, education or professional reputation cannot be evaluated solely by throughput and accuracy. Their design must consider explanation, contestability, privacy, proportionality and the lived experience of those affected.

Dignity also changes the language of transformation. Employees are not 'legacy costs' to be removed or 'human bottlenecks' to be bypassed. Such metaphors narrow managerial imagination and normalize disregard. A humanistic vocabulary does not prevent difficult decisions; it prevents leaders from hiding their human significance.

5.2 Purpose beyond efficiency

Efficiency is an important organizational value, but it is not a sufficient purpose. A hospital can process patients efficiently while failing to care; a university can produce credentials efficiently while weakening learning; and a company can personalize consumption efficiently while increasing manipulation. Humanistic leadership begins by asking what human good the organization exists to serve.

Purpose acts as a boundary condition for AI. It determines which problems should be automated, which activities should remain relational and which innovations should not be pursued even if they are profitable. Purpose is credible only when reflected in budgets, incentives, product choices and the distribution of benefits.

5.3 Freedom, agency and participation

Human flourishing requires agency: the experience of influencing one's work and acting for reasons one can endorse. Poorly implemented AI can undermine agency by prescribing actions, narrowing discretion and making employees accountable for outputs they did not meaningfully shape. Conversely, well-designed AI can expand agency by increasing access to expertise, reducing administrative burden and supporting creativity.

The leadership task is to design the boundary between recommendation and decision. High-impact decisions require meaningful human involvement, not merely a symbolic click. Participation in design and review helps ensure that technology strengthens rather than empties professional responsibility.

5.4 Relationships and the social nature of work

Work is a relational practice. Trust, belonging, care and recognition are not inefficiencies; they are conditions for cooperation and resilience. AI-mediated work can connect people, but it can also replace dialogue with dashboards and feedback with automated scores. Humanistic leaders protect spaces where people encounter one another as whole persons.

The test is not whether a process can be digitized, but what is lost when it is. Some interactions carry moral and emotional weight that requires presence, empathy and reciprocal understanding. Leadership must preserve these moments while using technology to support, not simulate, genuine relationship.

5.5 Flourishing as a multidimensional outcome

Humanistic leadership evaluates success through a broader account of value. Economic vitality matters because organizations need resources to endure. Yet flourishing also includes health, learning, meaning, fairness, connection, autonomy and contribution. These outcomes reinforce one another over time. A workforce that is exhausted, fearful or deskilled may produce short-term gains while weakening long-term adaptive capacity.

AI should therefore be assessed through a balanced value architecture. Productivity improvements must be considered alongside workload quality, distributive effects, trust, capability growth and environmental consequences. What is measured signals what leadership truly values.

6. THE HLAI GOVERNANCE ARCHITECTURE

The HLAI governance architecture comprises eight analytically distinct but interdependent domains. The dimensions are not encoded in an artificial acronym; they are named directly to preserve conceptual clarity and facilitate empirical operationalization. Together they specify what leaders must protect, decide, design and learn when AI becomes embedded in organizational processes.

Dimension	Core meaning	Leadership practices	Diagnostic question
Human Dignity	Treat every affected person as an end, never merely as a data source, risk score or productivity unit.	Conduct dignity impact assessments; protect privacy; provide understandable explanations; create appeal and remedy channels.	Would we consider this process acceptable if we were the person most vulnerable to its error?
Purpose Alignment	Anchor AI investment in a clearly articulated human and social purpose.	Define intended human benefit; identify prohibited uses; connect AI portfolio choices to mission and stakeholder value.	What worthwhile human outcome is this system meant to advance?
Moral Agency	Keep responsibility identifiable and preserve meaningful human choice.	Assign accountable owners; define escalation rights; prevent rubber-stamp oversight; train ethical reasoning.	Who can say no, who must explain, and who carries responsibility when the system causes harm?
Augmented Judgment	Use AI to strengthen, not replace, contextual and professional judgment.	Calibrate reliance; compare human and model reasoning; require uncertainty communication; design override protocols.	What does the system know, what does it not know, and what must a human understand before acting?
Relational Integrity	Protect trust, empathy, belonging and psychological safety during technological change.	Use participatory design; preserve high-empathy interactions; create safe reporting channels; monitor relational climate.	Does this technology help people work with one another or merely around one another?
Accountable Governance	Build transparent, multidisciplinary and continuous oversight across the AI lifecycle.	Maintain inventories; classify risk; audit data and performance; review vendors; document decisions; retire obsolete systems.	Can the organization reconstruct why this system was approved and how it has behaved over time?

Inclusive Value Distribution	Distribute benefits and burdens fairly across employees, customers, communities and future generations.	Measure differential impacts; redesign jobs; share productivity gains; ensure accessibility; consider environmental cost.	Who gains capability, who loses opportunity and who is absent from the decision?
Adaptive Learning	Create a culture that learns continuously from evidence, dissent, incidents and changing contexts.	Develop AI literacy; run pilots; use incident reviews; reward responsible challenge; update models and policies.	How will we know when our assumptions are wrong, and how quickly can we correct course?

Human Dignity

Treat every affected person as an end, never merely as a data source, risk score or productivity unit.

Dignity becomes operational when consequential systems preserve privacy, intelligibility, contestability and access to remedy. A dignity impact review should identify who is most exposed to classification error, exclusion or loss of voice, and whether safeguards are accessible in practice rather than only documented.

Human dignity is violated not only through dramatic abuse but through ordinary design choices that deprive people of context, voice or remedy. A dignitarian approach examines the entire experience: how data are collected, how categories are defined, how decisions are communicated and whether the person can challenge an outcome without fear or disproportionate effort.

Purpose Alignment

Anchor AI investment in a clearly articulated human and social purpose.

Purpose alignment requires a traceable link between an AI initiative, a legitimate stakeholder need and the organization's mission. Portfolio review should identify prohibited uses and reject applications whose commercial value depends on manipulation, excessive surveillance or avoidable harm.

Purpose is especially important because AI makes many things possible before organizations have decided whether they are desirable. A purpose-led organization creates explicit red lines and refuses uses that conflict with its identity, even when competitors adopt them. This self-limitation is not technological conservatism; it is strategic coherence.

Moral Agency

Keep responsibility identifiable and preserve meaningful human choice.

Moral agency is preserved when accountability is assigned to identifiable decision owners who possess authority, information and the right to halt deployment. Escalation routes must prevent both responsibility diffusion and symbolic human oversight.

Moral agency disappears when every participant assumes that someone else is responsible: the model, the vendor, the data, the policy or the executive committee. Leaders must prevent this diffusion. Accountability should be matched with authority, competence and access to information; otherwise it becomes blame without control.

Augmented Judgment

Use AI to strengthen, not replace, contextual and professional judgment.

Augmented judgment requires calibrated reliance rather than automatic deference. Teams should compare model and expert reasoning, communicate uncertainty, document overrides and examine cases in which contextual knowledge legitimately outweighs predictive output.

Augmentation is successful when the combined human-machine system performs better not only statistically but contextually. This requires attention to automation bias, overconfidence and deskilling. Professionals need practice in disagreeing with AI, documenting why and learning from cases where either human or machine judgment was superior.

Relational Integrity

Protect trust, empathy, belonging and psychological safety during technological change.

Relational integrity concerns the effects of AI on trust, recognition and psychological safety. Participatory design, transparent communication and protected reporting channels are especially important when implementation changes identity, status or employment expectations.

Trust cannot be automated. It grows through consistency, competence, benevolence and honest communication. During AI transformation, leaders should speak plainly about uncertainty, job implications and trade-offs. False reassurance creates deeper distrust than respectful candour.

Accountable Governance

Build transparent, multidisciplinary and continuous oversight across the AI lifecycle.

Accountable governance spans the full lifecycle: system inventory, risk classification, data and model documentation, vendor assurance, incident response, post-deployment monitoring and retirement. Control intensity should be proportionate to decision stakes and reversibility.

Governance must be proportionate to risk and continuous through the system lifecycle. Approval is not the end of responsibility. Data drift, model changes, altered user behaviour and new legal expectations can transform a once-acceptable system into a harmful one. Governance therefore combines documentation with active observation.

Inclusive Value Distribution

Distribute benefits and burdens fairly across employees, customers, communities and future generations.

Inclusive value distribution assesses who acquires capability, who bears transition costs and whether productivity gains improve job quality and access. Aggregate efficiency should not conceal concentrated disadvantage, accessibility barriers or environmental externalities.

Inclusive flourishing rejects the assumption that aggregate benefit justifies concentrated harm. Distribution matters. Leaders should examine whether productivity gains result in better work, learning and shared prosperity or simply intensify workloads and increase asymmetry between those who control technology and those who are controlled through it.

Adaptive Learning

Create a culture that learns continuously from evidence, dissent, incidents and changing contexts.

Adaptive learning treats incidents, near misses, overrides and stakeholder complaints as evidence for system revision. AI literacy, periodic review and protected dissent allow organizations to update assumptions before technical or social failures become institutionalized.

Intelligent learning depends on institutional humility. Organizations should expect surprises and design for correction. Near misses, employee concerns and stakeholder complaints are not public-relations problems to suppress; they are information about system quality. Learning organizations turn discomfort into prevention.

6.1 The HLAI Framework as a multilevel causal architecture

The Humanistic Leadership in Artificial Intelligence (HLAI) Framework positions leadership as the mediating governance layer between technological capability and organizational consequence. AI capabilities do not determine outcomes by themselves. Their effects depend on purpose, power, process design, human oversight, incentives, voice and learning. The framework therefore rejects both technological determinism and purely symbolic ethics. It proposes that humanistic leadership shapes three families of outcomes: human outcomes such as agency, trust, meaning, well-being and capability; organizational outcomes such as decision quality, innovation, resilience and legitimacy; and societal outcomes such as fairness, rights protection, inclusion and sustainable value.




Figure 1. Humanistic Leadership in Artificial Intelligence (HLAI) Framework.

Figure 1 illustrates the conceptual multilevel architecture of the Humanistic Leadership in Artificial Intelligence (HLAI) Framework. The model proposes that organizational AI capability—including technological infrastructure, data governance, digital competencies and AI readiness—forms the foundational layer enabling responsible AI adoption. These capabilities influence the HLAI governance architecture, comprising Human Dignity, Purpose Alignment, Moral Agency, Augmented Judgment, Relational Integrity, Accountable Governance, Inclusive Value Distribution and Adaptive Learning. These dimensions produce human, organizational and societal outcomes, while a continuous feedback loop uses stakeholder experience, incidents and governance evidence to refine leadership practices over time.

Textual representation:

Organizational AI Capability



HLAI Governance Layer



Human Outcomes + Organizational Outcomes + Societal Outcomes



Organizational Flourishing



Continuous Learning and Feedback

The downward logic of the model is complemented by feedback. Outcome data, stakeholder experience, incidents and near misses should revise governance arrangements and, when necessary, constrain or terminate an AI application. In this sense, humanistic AI leadership is not a one-time ethical approval but a continuous institutional learning process.

6.2 Comparative positioning of humanistic AI leadership

Dimension	Traditional leadership	Digital leadership	Efficiency-centred AI leadership	Responsible AI governance	Humanistic AI leadership	Distinctive HLAI contribution
Primary objective	Control and execution	Digital adaptation	Automation and productivity	Risk and compliance	Flourishing with performance	Integrates dignity, agency,

Dimension	Traditional leadership	Digital leadership	Efficiency-centred AI leadership	Responsible AI governance	Humanistic AI leadership	Distinctive HLAI contribution
						purpose and shared value
View of people	Resources and role holders	Users and change recipients	Data sources and system operators	Affected stakeholders	Persons with intrinsic worth and voice	Makes dignity a design and governance criterion
Decision logic	Managerial authority	Data-informed agility	Model optimization	Risk-based oversight	Augmented judgment and moral responsibility	Combines evidence, context, values and contestability
Success metric	Output and efficiency	Adoption and speed	Cost, accuracy and scale	Compliance, robustness and fairness	Human, organizational and societal value	Uses a balanced HAILI measurement architecture

Table 2. Comparative leadership paradigms in AI-enabled organizations.

7. LEADERSHIP TENSIONS AND HUMANISTIC RESPONSES

Tension	Humanistic interpretation	Leadership response
Speed vs. deliberation	AI rewards rapid experimentation, while high-impact decisions require reflection.	Use tiered governance: low-risk experimentation can move quickly; high-risk applications require deeper review, stakeholder participation and documented approval.
Personalization vs. privacy	Better predictions often depend on more personal data.	Apply data minimization, purpose limitation and proportionality. Ask whether the same benefit can be achieved with less intrusive data.
Consistency vs. context	Standardized models can reduce arbitrary variation but miss exceptional circumstances.	Create structured discretion and escalation pathways. Track both consistency gains and cases where context legitimately overrides the model.
Efficiency vs. meaning	Automation can remove effort but also remove learning, craftsmanship and human contact.	Redesign roles around higher-value contribution. Preserve developmental tasks and relational moments rather than automating everything technically feasible.
Prediction vs. self-fulfilling classification	Scores can shape opportunities and cause people to become what the model predicts.	Limit deterministic labels, communicate uncertainty and ensure periodic reconsideration. Avoid using historical disadvantage as an unquestioned forecast of future potential.

Transparency vs. complexity	Some models are difficult to explain fully.	Provide explanations appropriate to the decision and audience, including factors, limitations, uncertainty and recourse, rather than pretending that technical disclosure alone creates understanding.
Control vs. trust	Monitoring tools can increase visibility while damaging autonomy and psychological safety.	Use the least intrusive form of measurement, involve employees in design and separate developmental analytics from punitive surveillance.
Innovation vs. responsibility	Fear of risk can stop beneficial innovation, while enthusiasm can normalize careless deployment.	Adopt responsible experimentation: bounded pilots, reversible decisions, predefined success and harm criteria, and independent challenge.

Speed vs. deliberation

AI rewards rapid experimentation, while high-impact decisions require reflection.

Use tiered governance: low-risk experimentation can move quickly; high-risk applications require deeper review, stakeholder participation and documented approval.

These tensions cannot be resolved through a universal rule. The governance requirement is to make the competing values explicit, assign decision rights, select proportionate safeguards and document the conditions under which one value may legitimately take priority.

Personalization vs. privacy

Better predictions often depend on more personal data.

Apply data minimization, purpose limitation and proportionality. Ask whether the same benefit can be achieved with less intrusive data.

Consistency vs. context

Standardized models can reduce arbitrary variation but miss exceptional circumstances.

Create structured discretion and escalation pathways. Track both consistency gains and cases where context legitimately overrides the model.

Efficiency vs. meaning

Automation can remove effort but also remove learning, craftsmanship and human contact.

Redesign roles around higher-value contribution. Preserve developmental tasks and relational moments rather than automating everything technically feasible.

Prediction vs. self-fulfilling classification

Scores can shape opportunities and cause people to become what the model predicts.

Limit deterministic labels, communicate uncertainty and ensure periodic reconsideration. Avoid using historical disadvantage as an unquestioned forecast of future potential.

Transparency vs. complexity

Some models are difficult to explain fully.

Provide explanations appropriate to the decision and audience, including factors, limitations, uncertainty and recourse, rather than pretending that technical disclosure alone creates understanding.

Control vs. trust

Monitoring tools can increase visibility while damaging autonomy and psychological safety.

Use the least intrusive form of measurement, involve employees in design and separate developmental analytics from punitive surveillance.

Innovation vs. responsibility

Fear of risk can stop beneficial innovation, while enthusiasm can normalize careless deployment.

Adopt responsible experimentation: bounded pilots, reversible decisions, predefined success and harm criteria, and independent challenge.

8. AN OPERATING ARCHITECTURE FOR HUMANISTIC AI LEADERSHIP

8.1 Strategy and portfolio governance

Every AI initiative should begin with a problem statement framed in human and organizational terms. Instead of asking, 'Where can we deploy AI?', leaders should ask, 'Which stakeholder problem deserves to be solved, and is AI an appropriate means?' This reduces technology-driven solutionism and improves investment discipline.

The AI portfolio should be classified by impact and risk. Applications involving safety, employment, legal rights, essential services or vulnerable populations require stronger evidence and governance than tools used for low-stakes drafting. Boards should receive a balanced view of value, risk, workforce impact and unresolved uncertainty.

8.2 Decision rights and accountability

Organizations need a clear map of decision rights across the AI lifecycle. Business owners define purpose and acceptable outcomes; technical teams design and test; domain experts validate contextual relevance; risk and legal functions challenge; affected users contribute experience; and executive sponsors remain accountable for deployment. Vendor responsibility does not replace organizational responsibility.

A named accountable executive should be able to explain why the system exists, what evidence supported approval, which limitations remain and how incidents will be handled. Employees need explicit authority to pause or escalate questionable outputs without retaliation.

8.3 Human-centred process and job design

Process redesign should identify not only which tasks AI can perform but which human capabilities the organization wishes to strengthen. Repetitive administration may be automated so that employees can devote more time to judgment, creativity and care. Yet this benefit appears only when workloads and performance expectations are redesigned; otherwise saved time is simply filled with more volume.

Job design should protect pathways for learning. Junior professionals often acquire expertise through routine cases. If AI absorbs those cases, organizations must create alternative forms of deliberate practice, mentoring and progressive responsibility. Humanistic transformation prevents efficiency today from creating capability shortages tomorrow.

8.4 Culture and psychological safety

A responsible AI culture makes it normal to question outputs, admit uncertainty and report unintended effects. Leaders reinforce this culture through their reactions. When an employee raises a concern, the response should be curiosity before judgment. When a pilot fails, the organization should examine assumptions and system conditions rather than search immediately for an individual culprit.

Psychological safety does not mean low standards. It enables higher standards because problems surface earlier. Teams should conduct pre-mortems, red-team exercises and after-action reviews. Dissent should be structured into governance rather than left to personal courage alone.

8.5 Capability development

AI literacy must extend beyond prompt technique. Leaders and employees need to understand data provenance, model limitations, hallucination, bias, uncertainty, privacy, intellectual property and security. Equally important are human capabilities: critical thinking, ethical reflection, communication, empathy, imagination and systems thinking.

The aim is not to turn every employee into a data scientist. It is to create role-appropriate competence. Executives need enough understanding to govern; professionals need enough to use and challenge; technical teams need enough domain and ethical literacy to appreciate human consequences.

8.6 Metrics and incentives

What organizations measure will shape how AI is used. If success is defined only as cost reduction or throughput, teams will rationally ignore dignity, trust and capability. A humanistic scorecard combines performance with quality, fairness, well-being, learning and stakeholder outcomes.

Incentives should reward responsible escalation, documentation and improvement, not merely rapid deployment. Leaders should be cautious about productivity metrics that intensify surveillance or encourage employees to conceal the time required for judgment and care.

9. THE HUMANISTIC AI SCORECARD

The following scorecard translates the framework into indicators. It is not a universal reporting standard; organizations should select measures proportionate to context and risk. Quantitative metrics should be combined with qualitative evidence because dignity, meaning and trust cannot be captured by a single number.

Domain	Illustrative indicators	Interpretive question
Human value	Time released from low-value tasks; service quality; accessibility; stakeholder-reported benefit	Are benefits experienced by people or merely inferred from operational data?
Dignity and rights	Appeal availability; explanation quality; privacy incidents; consent quality; harmful use complaints	Can affected persons understand and challenge consequential outcomes?
Fairness and inclusion	Performance across groups; accessibility; distribution of gains and burdens; participation diversity	Does aggregate improvement conceal concentrated disadvantage?
Human agency	Override rate and quality; employee discretion; reported automation pressure; decision clarity	Are humans meaningfully deciding or merely endorsing machine output?
Trust and relationships	Psychological safety; stakeholder trust; perceived surveillance; quality of manager-employee dialogue	Has technology strengthened or weakened relational confidence?
Capability and learning	AI literacy; reskilling completion; learning transfer; incident closure; model improvement cycle time	Are people becoming more capable or more dependent?
Performance and resilience	Accuracy; productivity; quality; failure recovery; drift detection; continuity	Does the combined system remain reliable under changing conditions?
Governance integrity	Inventory completeness; risk reviews; audit findings; vendor transparency; retirement decisions	Can leaders' evidence responsible stewardship across the lifecycle?

9.1 Humanistic AI Leadership Index (HAILI)

To make the framework empirically researchable, this article proposes the Humanistic AI Leadership Index (HAILI). HAILI is not presented as a validated psychometric scale; it is a measurement architecture for future development. It contains eight domains aligned with the HLAI governance architecture and can be assessed through multiple sources: employee surveys, affected-stakeholder surveys, governance records, system logs, incident reviews and independent audits. A mature implementation should avoid self-report alone and should examine the gap between declared principles and experienced practice.

9.2 Proposed HAILI domains and indicators

HAILI domain	Illustrative indicators	Evidence source	Risk signal	Desired direction
Human dignity	Perceived respect; privacy protection; absence of demeaning classification	Survey, privacy audit, complaints	People feel reduced to scores or categories	Higher respect and lower dignity-related harm
Purpose alignment	Clarity of human benefit; prohibited-use boundaries; stakeholder value	Portfolio records, interviews	AI projects lack a defensible human purpose	Stronger purpose alignment

HAILI domain	Illustrative indicators	Evidence source	Risk signal	Desired direction
Moral agency	Named accountability; escalation use; refusal without retaliation	Governance logs, survey	No one can stop or own a harmful decision	Clearer responsibility and voice
Augmented judgment	Calibration, override quality, uncertainty communication	Decision logs, quality review	Automation bias or habitual rubber-stamping	Better calibrated human-machine judgment
Relational integrity	Psychological safety; trust; quality of consultation	Team survey, interviews	Fear, silence and deteriorating collaboration	Higher trust and speaking-up
Accountable governance	Risk tiering; monitoring; auditability; incident closure	Audit, incident data	Controls exist on paper but not in use	Higher governance effectiveness
Inclusive value distribution	Distribution of gains; accessibility; group-level performance	Impact assessment	Benefits concentrate while burdens shift to vulnerable groups	More equitable outcomes
Adaptive learning	AI literacy; reskilling; lessons from incidents; model retirement	Learning records, reviews	Capability decay and repeated failures	Continuous learning and timely adaptation

Table 5. Proposed domains and evidence architecture for the Humanistic AI Leadership Index (HAILI).

A future scale-development programme should begin with expert content validation and cognitive interviews, followed by exploratory and confirmatory factor analysis. Multilevel validation should test whether individual perceptions aggregate meaningfully to team or organizational levels. Criterion validity can be examined against trust, psychological safety, decision quality, innovation, well-being, turnover intention, incident frequency and stakeholder legitimacy. Cross-cultural invariance is essential because dignity, authority and voice may be enacted differently across institutional contexts.

10. IMPLEMENTATION ROADMAP: FROM PRINCIPLES TO PRACTICE

Phase 1 – Listen and map

Create an inventory of AI systems and shadow uses. Interview employees and affected stakeholders. Map decisions, data flows, risks, pain points and current governance gaps. Establish a baseline for trust, workload, capability and well-being.

Across all phases, the minimum governance record should state the intended benefit, evidence considered, unresolved uncertainty, accountable owner, affected stakeholders, escalation route and next review point.

Phase 2 – Define purpose and boundaries

Articulate the human purpose of AI adoption. Agree prohibited uses, high-risk categories and non-negotiable principles. Clarify the role of human judgment and the rights of affected persons.

Phase 3 – Design governance and participation

Assign owners and decision rights. Create multidisciplinary review. Include employees, customers or citizens where relevant. Develop documentation, escalation, appeal and incident processes.

Phase 4 – Pilot responsibly

Run bounded, reversible pilots. Define value, harm and learning criteria before launch. Compare human-only, AI-only and combined performance where appropriate. Observe workflow and relational effects, not just technical accuracy.

Phase 5 – Redesign work and build capability

Remove low-value tasks, redesign workloads, preserve learning pathways and provide role-specific AI literacy. Support identity transitions and explain how contribution will be recognized.

Phase 6 – Scale with evidence

Scale only when evidence supports value and acceptable risk. Monitor drift, group impacts, behaviour changes and unintended consequences. Maintain human review proportionate to stakes.

Phase 7 – Learn, report and renew

Publish internal performance and incident insights. Review systems periodically, including the possibility of retirement. Update policies as technology, law and stakeholder expectations evolve.

11. SECTOR APPLICATIONS

11.1 Technology and software engineering

In software development, generative AI can accelerate coding, testing, documentation and incident analysis. Humanistic leadership requires attention to security, intellectual property, reliability and the development of junior engineers. Teams should review generated code, track provenance, protect confidential data and ensure that speed does not displace architectural understanding.

Engineering leaders should treat maintainability and human comprehension as quality attributes. Code that works but cannot be responsibly understood, audited or supported may create future fragility.

11.2 Human resources and people analytics

AI can support recruitment, workforce planning and learning personalization, but employment decisions directly affect dignity and opportunity. Historical data may encode past exclusion; behavioural proxies may be weak or intrusive; and automated communication can make candidates feel invisible.

Humanistic practice requires job-relevant validation, bias testing, accessible explanations, human review and appeal. People analytics should support development and organizational insight rather than covert surveillance or deterministic labeling.

11.3 Healthcare

AI can improve diagnosis, triage and administrative efficiency, but clinical care contains uncertainty, vulnerability and moral responsibility. Patients need to know when AI materially influences care, clinicians need authority to disagree, and systems must be evaluated across populations and contexts.

The humanistic objective is not to replace the therapeutic relationship. It is to provide clinicians with better support and more time for listening, explanation and shared decision-making.

11.4 Education

Generative AI can expand tutoring, feedback and access to knowledge. It can also weaken authorship, critical thinking and the teacher-student relationship when used without pedagogical purpose. Humanistic leadership in education asks what learners should become, not merely what outputs they should produce.

Institutions should redesign assessment, teach AI literacy and preserve intellectual struggle where it is necessary for growth. The goal is augmented learning, not frictionless completion.

11.5 Public administration and essential services

Public-sector AI carries special obligations because citizens may have limited choice and decisions can affect rights or essential support. Transparency, due process, accessibility and democratic accountability are central. Efficiency cannot justify systems that make public power harder to understand or challenge.

Humanistic public leadership should prioritize service quality, equality before the law and the ability of citizens to receive a reasoned human response when automated processes fail.

11.6 Implications for engineering and technology management

For engineering and technology management, the HLAI Framework shifts responsible AI from a specialist ethics activity to an operating responsibility shared across the system lifecycle. Product managers should define human benefit and prohibited uses before feature development. Data and machine-learning teams should document data provenance, uncertainty, performance variation and foreseeable misuse. Systems engineers should design escalation, fallback and graceful degradation. Cybersecurity and safety teams should integrate model-specific risks into existing assurance processes. Human resources should anticipate role transition, capability needs and distributive effects. Senior leaders should make explicit who may pause, override or retire a system.

The framework is especially relevant to Industry 5.0 and smart manufacturing, where optimization can improve productivity while simultaneously intensifying surveillance, narrowing worker discretion or transferring risk to frontline employees. A humanistic design asks not only whether a system improves throughput, but whether it strengthens worker competence, reduces avoidable danger, preserves meaningful control and distributes gains fairly. In agile software environments, humanistic governance should be embedded into product backlogs,

definition-of-done criteria, architecture reviews and post-deployment monitoring rather than added at the end as an ethics checklist.

ISO/IEC 42001, ISO/IEC 23894, the NIST AI RMF and the EU AI Act provide complementary governance foundations. The distinctive contribution of HLAI is to connect those structures to leadership behaviour, employee experience and organizational purpose. Compliance can specify minimum controls; leadership determines whether those controls become living practices, whether inconvenient evidence is heard and whether commercial pressure overrides declared principles.

12. PROPOSITIONS FOR FUTURE EMPIRICAL RESEARCH

P1. Higher HLAI maturity will be positively associated with employee trust in organizational AI adoption.

P2. The positive relationship between AI capability and organizational performance will be stronger when meaningful human judgment is preserved.

P3. Psychological safety will mediate the relationship between accountable AI governance and early incident reporting.

P4. Participatory AI design will reduce perceived surveillance and increase implementation quality, particularly in high-discretion work.

P5. AI-enabled productivity gains will improve sustainable performance only when accompanied by job redesign and capability development.

P6. Explanation, appeal and remedy mechanisms will increase perceived procedural fairness and stakeholder legitimacy in consequential AI decisions.

P7. Organizations using multidimensional human-impact measures will identify unintended consequences earlier than organizations relying predominantly on efficiency metrics.

P8. The effects of HLAI maturity on trust and performance will vary according to national culture, regulatory context, task interdependence and decision stakes.

These propositions can be tested through longitudinal field studies, multilevel surveys, comparative case research, experiments and mixed-method evaluations. Future research should develop and validate measures for each HLAI dimension, examine interactions among dimensions and investigate whether the framework predicts trust, innovation, well-being, fairness and long-term performance.

P9. HLAI maturity will strengthen the relationship between AI capability and organizational performance through improved adoption quality and error detection.

P10. HLAI maturity will weaken the relationship between employee AI exposure and loss of meaningful work through job redesign, learning opportunities and perceived contribution.

P11. The effect of explainability on trust will be mediated by contestability and perceived procedural fairness.

P12. The relationship between responsible AI controls and stakeholder legitimacy will be strengthened by consistency between leadership statements, incentives and operational decisions.

13. DISCUSSION

The argument developed in this article challenges two common simplifications. The first is technological determinism: the belief that AI will produce a single inevitable future of work. In reality, outcomes depend on choices about ownership, purpose, job design, governance and distribution. The second is ethical add-on thinking: the idea that values can be addressed after the technical system is built. Values are already present in problem formulation, data selection, optimization targets and interface design. Humanistic leadership must therefore enter at the beginning.

The HLAI Framework is differentiated from adjacent approaches by its unit of analysis and its integration mechanism. NIST AI RMF and ISO/IEC 42001 primarily organize risk-management and management-system controls; Shneiderman's human-centred AI emphasizes reliable, safe and trustworthy system design; humanistic management establishes dignity and flourishing as normative ends; and Industry 5.0 links technology to human-centricity, resilience and sustainability. HLAI connects these contributions at the level of leadership practice. It specifies how purpose, decision rights, employee voice, job design, contestability, incentives and organizational learning mediate the relationship between technical capability and experienced outcomes. HLAI then translates those mechanisms into a provisional measurement architecture. The novelty therefore lies not in proposing another list of ethical principles, but in linking normative aims, managerial mechanisms and evaluable evidence within one multilevel model.

The framework also redefines competitive advantage. In an environment where access to AI becomes widespread, differentiation will depend increasingly on organizational judgment, trust, domain knowledge and the quality of human collaboration. A culture capable of questioning technology responsibly may innovate faster because it detects problems earlier and learns more openly. Humanism and performance are not opposing agendas; under conditions of uncertainty, they can become mutually reinforcing.

Nevertheless, humanistic language can be misused. Organizations may celebrate people while deploying systems that intensify control and insecurity. The framework therefore insists on evidence. Humanistic leadership must be visible in decision rights, resource allocation, metrics, appeals, job quality and the distribution of benefits. Without institutional translation, values remain branding.

A further implication concerns the role of leaders themselves. AI can generate analyses, speeches and strategies that appear authoritative. Leaders may be tempted to outsource not only tasks but reflection. Yet leadership is not the production of polished text. It is the willingness to enter reality, listen to those affected, make accountable choices and remain present for consequences. In this sense, AI does not make leadership less human. It exposes whether leadership was human in the first place.

14. MANAGERIAL IMPLICATIONS

For boards, the principal implication is that AI oversight belongs within strategy and corporate governance, not only cybersecurity or compliance. Boards should ask how AI changes the organization's purpose, risk profile, workforce, stakeholder relationships and sources of legitimacy.

For chief executives, the challenge is coherence. Public commitments to human-centred AI must match operational incentives. Leaders should resist announcing scale targets before establishing evidence and governance. They should communicate honestly about job impact and share gains through better work, development and opportunity.

For technology leaders, the framework expands the definition of quality. Accuracy, latency and cost remain important, but quality also includes explainability, contestability, usability, accessibility, security and social consequence. Technical excellence becomes inseparable from responsible integration.

For human resource leaders, AI transformation should be treated as an organizational development process. Skills, roles, identity, fairness, trust and leadership behaviour require active design. Reskilling promises must be translated into protected time, accessible pathways and real opportunities.

For line managers, humanistic leadership becomes highly practical: explain why a tool is being used; invite challenge; review outputs thoughtfully; preserve dialogue; notice workload and anxiety; and defend the right to escalate. Culture is built in these daily interactions.

14.1 Leadership capability model for the AI transition

Humanistic AI governance depends on capabilities that cannot be delegated to a compliance office. Leaders at every level need enough technological literacy to ask informed questions, enough ethical literacy to recognize value conflicts, enough relational skill to hear fear and dissent, and enough strategic discipline to connect AI investments to a defensible human purpose. The required profile is not that of the leader as engineer, philosopher or therapist, but of an integrator who can convene these perspectives and make accountable choices under uncertainty.

Five capability clusters are particularly important. First, sensemaking enables leaders to interpret probabilistic outputs, distinguish evidence from confidence theatre and recognize when context invalidates a model recommendation. Second, moral imagination enables them to anticipate how a system may affect people who are absent from the design room, especially those with less power. Third, dialogic leadership creates conditions in which engineers, employees, users and external stakeholders can challenge assumptions without fear. Fourth, governance craftsmanship converts principles into decision rights, escalation routes, monitoring and consequences. Fifth, regenerative learning treats incidents, overrides and stakeholder criticism as sources of institutional improvement rather than threats to reputation.

Capability cluster	Observable leadership behaviour	Failure mode when absent	Development method
AI sensemaking	Questions uncertainty, limitations, data quality and contextual fit	Automation bias or rejection based on misunderstanding	Scenario analysis, model cards, decision reviews
Moral imagination	Anticipates effects on vulnerable and absent stakeholders	Narrow optimization and hidden externalities	Dignity impact exercises, stakeholder panels
Dialogic leadership	Invites dissent and protects speaking-up	Silence, groupthink and symbolic consultation	Psychological-safety routines, red-team dialogue

Governance craftsmanship	Defines owners, thresholds, escalation and retirement rules	Principles without operational force	Lifecycle governance simulations, accountability mapping
Regenerative learning	Uses incidents and overrides to improve systems and work design	Repeated failure and defensive concealment	After-action reviews, learning dashboards, cross-functional retrospectives

Table 6. Leadership capabilities for humanistic AI transformation.

Capability development should be tied to real decisions rather than delivered only as abstract training. Executive teams can examine a live AI portfolio and practise stopping or redesigning a project whose benefits do not justify its human costs. Managers can review cases in which model confidence conflicts with professional judgment. Technical teams can participate in stakeholder conversations and job-shadowing to understand lived consequences. Boards can test whether they receive information on human impact with the same seriousness as financial and cybersecurity risk. These practices turn humanistic leadership from an aspiration into a repeatable organizational competence.

This capability model also has implications for leadership selection and evaluation. Organizations should not reward leaders solely for adoption speed or savings generated by automation. Evaluation should include the quality of participation, the credibility of safeguards, capability growth, distribution of benefits, incident learning and the preservation of meaningful work. When incentives recognize only short-term efficiency, even well-designed governance systems are likely to be bypassed. Humanistic AI leadership therefore requires alignment among competence, authority, information and reward.

15. LIMITATIONS

This study is conceptual and does not empirically validate the proposed framework. The dimensions may overlap and their relative importance may differ across sectors, cultures and risk levels. Future studies should test discriminant validity and examine whether a shorter or context-specific structure is more useful.

The article synthesizes a rapidly changing field. Technological capabilities, regulation and public expectations will continue to evolve. The framework is designed around enduring humanistic principles, but its practices and indicators must be reviewed regularly.

Finally, the positive potential of humanistic leadership should not obscure structural questions of market power, labour bargaining, environmental cost and geopolitical inequality. Organizational leadership is necessary but not sufficient. Responsible AI also requires effective law, professional standards, public institutions and societal debate.

16. CONCLUSION

Artificial intelligence is a mirror and an amplifier. It reveals what organizations value and increases their capacity to act on those values. Where people are treated as costs, AI can intensify extraction. Where authority is opaque, AI can deepen distance. But where dignity, purpose and shared flourishing guide leadership, AI can release human energy from repetitive work, expand access to knowledge and strengthen the quality of judgment.

Humanistic leadership in the age of AI is therefore not nostalgia for a pre-digital workplace. It is a forward-looking governance capability. It combines moral agency with technological literacy, empathy with evidence, participation with accountability and innovation with restraint. The HLA governance architecture offers a structured way to make this combination operational.

The decisive question for leaders is not whether machines will become more intelligent. They will. The decisive question is whether organizations will become wiser: more capable of choosing worthy purposes, recognizing human vulnerability, distributing benefits fairly and learning from consequences. The future of work will not be secured by technology alone. It will be shaped by leaders who remember that progress becomes meaningful only when it helps people and communities to flourish.

The decisive question of the AI age is not whether machines will become more capable. They will. The decisive question is whether organizations will become more humane as their capabilities expand. Humanistic leadership rejects the false choice between innovation and dignity. It proposes a more demanding standard: innovation that deserves trust because it enlarges human agency, protects the vulnerable, improves judgment and creates shared value. The future of leadership will therefore be measured not by how much intelligence an organization can automate, but by how wisely it governs power, how courageously it accepts responsibility and how deeply it enables people and communities to flourish.

DECLARATIONS

Originality statement: This manuscript is an original conceptual article prepared for submission to the Excel International Journal of Technology, Engineering & Management. It has not been published elsewhere and does not report fabricated empirical data.

Conflict of interest: The author declares no conflict of interest related to this manuscript.

Funding: No external funding was received for this study.

Data availability: No new dataset was created or analysed because the study is conceptual and review-based.

Ethics approval: Not applicable.

APPENDIX A. HUMANISTIC AI LEADERSHIP AUDIT

This practical audit can be used before approval, during periodic review or after a significant incident. A “no” answer does not automatically prohibit deployment, but it requires documented justification, mitigation and accountable ownership.

1. Purpose: Is the human or social problem clearly defined before the technology is selected?
2. Necessity: Is AI genuinely needed, or could the objective be achieved through a simpler and less risky process?
3. Dignity: Have we considered the experience of the person most vulnerable to an erroneous outcome?
4. Data: Are data relevant, lawful, proportionate, sufficiently representative and protected throughout their lifecycle?
5. Agency: Does the process preserve meaningful human choice rather than a symbolic approval click?
6. Accountability: Is there a named owner with authority, competence and resources to act?
7. Explanation: Can affected people understand the principal reasons, limitations and consequences of the output?
8. Appeal: Is there an accessible route to challenge, correct and remedy a consequential decision?
9. Fairness: Have performance and errors been examined across relevant groups and contexts?
10. Participation: Have employees, users or communities affected by the system contributed to its design or review?
11. Work quality: Will the system improve the content of work rather than merely increase volume and surveillance?
12. Capability: Are learning pathways protected, especially for junior employees whose routine tasks may be automated?
13. Relationships: Are high-empathy or trust-critical interactions being preserved for genuine human engagement?
14. Psychological safety: Can people question or stop the system without fear of retaliation?
15. Uncertainty: Are model confidence, data limitations and unknowns communicated honestly?
16. Override: Are criteria for human override clear, monitored and learned from?
17. Security: Have misuse, leakage, adversarial manipulation and dependency risks been tested?
18. Vendors: Are third-party responsibilities, data practices, update rights and audit access contractually clear?
19. Environmental impact: Has the energy and resource cost been considered in relation to the value created?
20. Distribution: Who receives the gains, who bears the burdens and how will productivity benefits be shared?
21. Metrics: Does the scorecard include dignity, trust, fairness, learning and well-being alongside efficiency?
22. Pilot design: Is the first deployment bounded, reversible and governed by predefined success and harm criteria?
23. Monitoring: Are drift, behaviour changes, incidents and stakeholder complaints reviewed continuously?
24. Retirement: Are there criteria and authority for suspending or retiring a system that no longer serves its purpose?
25. Leadership integrity: Would senior leaders be willing to explain and defend the system publicly to those most affected?

REFERENCES

- Amershi, S., Weld, D., Vorvoreanu, M., et al. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-13.
- Ashforth, B. E., & Kreiner, G. E. (1999). “How can you do it?” Dirty work and the challenge of constructing a positive identity. *Academy of Management Review*, 24(3), 413-434.
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30.
- Binns, R. (2018). Fairness in machine learning: Lessons from political philosophy. *Proceedings of Machine Learning Research*, 81, 149-159.
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work*. NBER Working Paper No. 31161.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Deci, E. L., Olafsen, A. H., & Ryan, R. M. (2017). Self-determination theory in work organizations. *Annual Review of Organizational Psychology and Organizational Behavior*, 4, 19-43.
- Dignum, V. (2019). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer.

- Edmondson, A. C. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350-383.
- European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.
- Floridi, L., Cowsls, J., Beltrametti, M., et al. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28, 689-707.
- Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Pitman.
- Grant, A. M. (2008). The significance of task significance. *Journal of Applied Psychology*, 93(1), 108-124.
- International Labour Organization. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality*. ILO.
- ISO/IEC. (2023). ISO/IEC 42001:2023 Information technology—Artificial intelligence—Management system. International Organization for Standardization.
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410.
- Microsoft & LinkedIn. (2024). *Work Trend Index Annual Report: AI at work is here. Now comes the hard part*.
- NIST. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology.
- OECD. (2019). *Recommendation of the Council on Artificial Intelligence*. Organisation for Economic Co-operation and Development.
- Pirson, M. (2017). *Humanistic management: Protecting dignity and promoting well-being*. Cambridge University Press.
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation-augmentation paradox. *Academy of Management Review*, 46(1), 192-210.
- Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68-78.
- Schein, E. H., & Schein, P. A. (2018). *Humble leadership: The power of relationships, openness, and trust*. Berrett-Koehler.
- Schwartz, B., & Sharpe, K. E. (2010). *Practical wisdom: The right way to do the right thing*. Riverhead Books.
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press.
- Spreitzer, G. M., Sutcliffe, K., Dutton, J., Sonenshein, S., & Grant, A. M. (2005). A socially embedded model of thriving at work. *Organization Science*, 16(5), 537-549.
- UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. United Nations Educational, Scientific and Cultural Organization.
- UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO.
- Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the unexpected: Sustained performance in a complex world* (3rd ed.). Wiley.
- World Health Organization. (2021). *Ethics and governance of artificial intelligence for health*. WHO.
- Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (NIST AI 600-1)*. National Institute of Standards and Technology.
- International Labour Organization. (2025). *Generative AI and jobs: A refined global index of occupational exposure*. ILO Working Paper 140.
- ISO/IEC. (2023). ISO/IEC 23894:2023 Information technology—Artificial intelligence—Guidance on risk management. International Organization for Standardization.
- ISO/IEC. (2025). ISO/IEC 42006:2025 Artificial intelligence—Requirements for bodies providing audit and certification of artificial intelligence management systems. International Organization for Standardization.
- OECD. (2024). *Recommendation of the Council on Artificial Intelligence (updated 2024)*. Organisation for Economic Co-operation and Development.
- UNESCO. (2024). *AI competency framework for teachers*. United Nations Educational, Scientific and Cultural Organization.

How do I cite this article?

Ignacio Bonasa Alzuria, Humanistic leadership in the age of artificial intelligence: A human-centred governance and measurement framework for ai-enabled organizations, *Excel International Journal of Technology, Engineering and Management*, 2026; Volume -13, Issue-3_Page_85-103.
<https://doi.org/10.64971/j.cph.eijtem.v13.i3.9.2026>



This is an open access article under the CC BY-NC-ND license
(<http://creativecommons.org/licenses/by-nc-nd/4.0/>)