

The Stakeholder Perspective in the Generative AI Era¹

A new series of stakeholder-centered short guidance articles on projects, people, and performance in the age of Generative Artificial Intelligence

STAKEHOLDERS IN THE AI ERA: WHO ARE THEY?²

by Massimo Pirozzi

Stakeholders, who are they? Stakeholders are the people, groups or organizations that are central to every project. They are the ones who do the work and who benefit from the project. Their satisfaction is as important to success as their contribution to progress. Therefore, all projects deliver outputs, outcomes, and value through human collaboration and reciprocal influence.

Yet, in just a few years, the context has changed profoundly. Generative AI (GAI) has transformed the conception, execution, and communication of projects, with hybrid ecosystems increasingly carrying out initiatives through continuous collaboration between people and "intelligent" systems. The distinction between the two has become increasingly blurred: artificial agents – large language models, recommendation algorithms, predictive engines, and autonomous decision-support platforms – can now influence strategies, communication patterns, and even outcomes in ways that were once exclusively human.

Can Artificial Intelligence Really Be a Stakeholder?

This evolution raises an essential question: can an AI system be considered a stakeholder if stakeholders are, by definition, human individuals or groups who have interests, who influence, or who are influenced by a project? This is not just an academic question; artificial intelligence systems already act as participants with a stake in outcomes by enabling communication, processing knowledge, supporting decision-making, and reflecting collective preferences. May generative AI be considered a stakeholder? It is

¹The contents of this Article are extracted and synthesized from Chapter 15, "Stakeholders in the AI Era – Who Are They?", of the Author's book "The Stakeholder Perspective: Projects, People and Performance in the Generative AI Era" (2nd Edition), Taylor & Francis / CRC Press, Boca Raton (FL), U.S.A., June 2026, ISBN 978-1-041-14837-1.

²How to cite this paper: Pirozzi, M. (2026), Stakeholders in the AI Era – Who Are They? The Stakeholder Perspective in the Generative AI Era; Article 1 in a new series of stakeholder-centered short guidance articles on projects, people, and performance in the age of Generative Artificial Intelligence, *PM World Journal*, Vol. XV, Issue IX, September.

not, because stakeholders are individuals and collectives, while GAI is non-human. Conversely, can we regard GAI as a stakeholder and thereby manage our relationship with it? Yes, because GAI emulates human behavior and engages with us in human language, and all stakeholder relationships are founded on behaviors (Pirozzi, 2024): AI stakeholders add a novel layer to stakeholder relationships without replacing the human ones.

In this article, the first in a new series on the Stakeholder Perspective in the Generative AI Era, we examine how the stakeholder landscape has evolved from the classic categories of Providers, Purchasers, Investors, and Influencers, to also include AI-powered entities that can reinforce or disrupt human behavior and expectations. The change is not just technological, but also about identity and relationships. The traditional approach to stakeholder identification was based on authority and interest, assuming that stakeholders act intentionally and are guided by values and motivations. This is no longer the case for systems without consciousness: AI systems do not have their own goals and interests and should therefore be considered an agency rather than an intelligence (Floridi, 2023). But their outputs influence plans, affect decisions, and transform perceptions, showing what can be called “functional agency” — an operational capability coming from computational processing rather than from human consciousness.

This functional perspective allows project managers to evaluate AI participation similarly to human participation – applying the same criteria of interest, influence, legitimacy, and impact – while keeping human responsibility and accountability as paramount. If stakeholders are defined by their interests, expectations, and behaviors, the first challenge in applying this concept to artificial entities is understanding what it means for an AI system to have an “interest.” An AI stakeholder’s “interest” is the operational goal built into its design or prompt: it tells the system what to do and changes its actions in a given situation. This requires a shift from psychological to functional intentionality (Floridi, 2019): instead of thinking of systems as beings with motivations, we consider them as actors with operational goals that affect the stakeholder ecosystem through a “chain of delegation” in which the designer’s objectives are embedded in the algorithm, the trainer’s biases are reflected in the training data, and the user’s intent is conveyed through prompts and operational parameters.

A Shifting Landscape of Power and Communication

Several interconnected dynamics are already reshaping the traditional balance of power among stakeholders:

- Data as currency. Those who own or manage datasets have a structural advantage, since their predictive analysis may influence decision-making and prioritization.

- Algorithmic mediation. AI often acts as a "mediator," allowing stakeholders who influence the algorithms' outputs to gain power that exceeds their formal role in the organization.
- Speed of analysis. Near-real-time AI-driven analysis may enable faster and more objective decisions, but excessive reliance on it may reduce accuracy, empathy, and time for reflection.
- "Democratized" expertise. Tools that once required specialist skills are now widely available, expanding analytical participation while challenging professional boundaries.
- Bias transformation. While AI may mitigate some forms of human bias, it introduces algorithmic bias (Bostrom, 2014) through data quality, prompt design, and modeling assumptions, creating new asymmetries of power.

A further transformation concerns communication itself. While traditional stakeholder communication was, by definition, communication among humans, AI stakeholders introduce new dynamics, and communication in the AI era becomes continuous and adaptive. Chatbots, recommendation engines, and predictive sentiment analysis provide additional possibilities for both evidence-based and constant interaction: stakeholders receive instant feedback, personalized content, and real-time updates, while decision-makers access predictive dashboards that signal opportunities and risks before they emerge. But this constant flow, while it makes for efficiency and responsiveness, also risks creating a simulation of authenticity. As interactions become more digital, it is harder to distinguish real understanding from algorithmic approximation, and stakeholders may be interacting with digital doubles – interfaces that mimic empathy or authority but possess no real agency. Trust is therefore moving from people to systems and from communication to design, which requires paying continuous attention both to human narratives and to computational outputs to ensure that automated mediation does not distort how people feel or what they mean.

This continuous flow of information also raises stakeholder expectations. As AI speeds up communication and analysis, stakeholders become used to faster results: what once counted as "responsive" behavior is now simply expected, and both internal and external stakeholders demand immediate responses, evidence-backed explanations, and direct contact. AI-driven dashboards, notifications, and predictive analytics make things more transparent than ever, but they also make stakeholders less patient, especially when things remain unclear or take longer than anticipated.

At the same time, the AI revolution expands the stakeholder network itself. Along with human participants, it brings in new artificial players and new human roles that come directly from the AI field – like data scientists, AI ethicists, platform providers, regulatory

agencies, and advocacy groups that focus on fairness, bias, or sustainability. Horizontally, every AI deployment connects the project's ecosystem with external domains – cloud infrastructures, open-source communities, and global data services – forming digital interdependencies that go beyond contractual boundaries. Because influence and accountability now extend beyond hierarchical and organizational borders, and because AI systems operate continuously and evolve with data, the stakeholder network, traditionally quite stable and predictable, becomes dynamic, behaving like a living organism that learns, adapts, and responds through digital feedback. Managing this kind of network requires a shift from mere control to orchestration: project managers must integrate cues derived from ethics, computational signals, and human insight, becoming leader-conductors able to harmonize human and artificial intelligences.

Influence, Archetypes, and Accountability

As AI systems increasingly influence decision-making, they can be usefully understood through the type of influence they exert, which can be indirect, direct, or informative (Pirozzi, 2025). *Indirect influence* is supportive or mediated: the AI agent analyzes data, makes suggestions, or lays out possible outcomes, but the final choice remains with a person. *Direct influence* is autonomous: the artificial agent operates and decides on its own, without real-time human supervision, as in high-frequency trading, recommendation engines, or self-driving conversational agents. *Informative influence* arises from AI's growing role in forming opinions and disseminating content, with the associated risk of “botshit” – plausible but false information that can mislead stakeholders and decision-makers.

Empirical observation across organizations identifies three recurring archetypes of AI-stakeholder behavior, each with distinct governance implications:

- The Amplifier. Improves consistency, speed, and effectiveness of stakeholder interactions. Risk: it strengthens biases or errors in data and instructions just as much as it strengthens strengths.
- The Interpreter. Translates complex information into narratives accessible to different audiences, bridging technical and managerial perspectives. Risk: simplification may distort meaning.
- The Autonomous Collaborator. Executes tasks or generates outputs with limited human involvement. Risk: ambiguous or unverified outputs complicate accountability.

This last point leads directly to the question of accountability. In traditional project environments, accountability is associated with specific roles – the sponsor, the project manager, the work package leader. When AI is added to ecosystems, these “borders of responsibility” become less evident, and the risk is the diffusion of responsibility: a

situation in which numerous actors contribute to an outcome, and no single individual controls the full chain of events that produced it. Three complementary principles help restore accountability in AI-mediated initiatives:

- **Traceability.** Every action or suggestion made by AI must be traceable back to the person or group who designed, configured, or approved the use of the system.
- **Auditability.** AI-related processes and conclusions must be examinable, with data sources, model parameters, and decision logic properly documented.
- **Attributability.** AI systems must always allow the identification of the last human responsible for the operational duty assigned to them.

These principles rest on a simple and unbreakable rule: automation cannot take over responsibility (Bryson, 2018); the persons and organizations that design and operate an algorithm are the only ones who can be held legally or ethically liable for it. The same logic applies to trust. In hybrid communication networks, stakeholders base their trust in AI-mediated information on three interrelated factors – technical reliability, organizational integrity, and relational authenticity – which together generate what can be called “synthetic trust”: a form of trust that is grounded in human values, but enabled by technology, and therefore fragile, since a single error or opaque outcome can quickly erode it.

Toward Stakeholder Perspective 4.0

In summary, the stakeholder ecosystem of the AI era is more complex, interconnected, and “intelligent” than it was before: algorithms are changing how people work, access to data is redistributing power, communication is accelerating, expectations are rising, and new stakeholders – some real, some virtual – are entering the picture. Nevertheless, relationships, trust, and a shared view of value remain the most critical determinants of stakeholder success, and are now best developed within interdependent hybrid networks where people and systems continuously learn from, and adapt to, each other.

On this basis, the Stakeholder Perspective can be extended into a more inclusive definition, here proposed as Stakeholder Perspective 4.0: “A stakeholder is any entity that impacts or is affected by the objectives, outcomes, or relationships of a project or organization, through intentional or functional participation, within an ecosystem grounded in human values, ethical accountability, and shared purpose.”

This definition respects the moral primacy of the human while recognizing the operational reality of the artificial: influence is no longer limited to individuals, but value remains, and will remain, a human construct, since only humans can understand what value means. Recognizing AI stakeholders is only the beginning of a broader transformation.

Before returning to the question of accountability, it is worth clarifying what an AI stakeholder actually is, since although AI is not a “classic” stakeholder – stakeholders are typically persons or groups of individuals, and AI is not human – we can consider AI as a virtual stakeholder: this allows us to manage our relations with it and engage with it, as AI behaves similarly to humans by interacting with us in human language, and all stakeholder relationships are founded on behaviors (Pirozzi, 2024). AI-Stakeholders therefore introduce a new dimension, rather than a new category, within the traditional map of stakeholder communities: they have specific characteristics because they are virtual and digital, and they may act as simulators and emulators of real stakeholders, remaining intrinsically neutral in terms of interests, until such time as we ask them to share our interests. In practice, project managers can prompt AI with a persona or role pattern, asking it to simulate the behaviors and languages of the four traditional stakeholder communities:

- As a Provider, AI can assist in managing resources, time, costs, risks, quality, contracts, control, issues, and communications.
- As a Purchaser, AI can assist in managing requirements, expectations, feedback, requests for change, acceptance criteria, value, outcomes and benefits, and communications.
- As an Investor, AI can assist in budget oversight, strategic alignment, strategic risk management, and communications, from reporting to institutional relations.
- As an Influencer, AI can assist in ensuring regulatory compliance, in public relations, and in managing communications with a large variety of communities.

Recognizing AI as a stakeholder does not equate it with human participants in moral or legal terms: AI has power but not responsibility, and its influence stems from human delegation, not from autonomy in the ethical sense. Humans remain the only moral agents in the system – the only ones capable of empathy, judgment, and accountability. Recognizing AI stakeholders is not meant to humanize machines, but to clarify the scope of human duty: to ensure that artificial participation reinforces, rather than undermines, human decision-making, creativity, and ethics.

Bias, Trust, and the Architecture of Governance

This growing sophistication of AI systems brings the discussion back, inevitably, to bias. When people contemplate the ethics of AI, bias is generally the first thing that comes to mind: AI systems learn from data, which reflects the imperfections of the environment that produced it, and bias has a significant impact on how stakeholders are managed. An AI system analyzing human emotions may struggle to interpret cultural expressions correctly; a predictive model for resource allocation might favor groups with larger or better-organized datasets; a generative assistant responsible for project communications

could replicate the biases present in its training data. This generates what can be called the bias paradox:

- AI promises fairness and objectivity, but it can only be neutral if the data and prompts it uses – which come from humans – are themselves neutral.
- The more we rely on AI to remove human bias, the more likely we are to make algorithmic bias a permanent, and less visible, part of our systems.

Resolving this paradox requires more than technological improvements such as bias-detection algorithms: it requires knowing how our relationships with other people change the way we think, i.e., project leaders should treat AI-generated analysis as an informed opinion rather than as an established fact, and cultivating a diverse group of people in charge – different genders, cultures, skills, and points of view – remains a key safeguard against algorithms that only “think” in one direction.

Bias directly affects the second pillar of every stakeholder relationship: trust. In relationships with stakeholders, trust has always been the hidden capital, which grows slowly when people are honest, open, and ready to communicate. In hybrid communication networks, people, groups, and algorithms all trust each other, and trust becomes “triangular”, resting on three interrelated factors: technical reliability – trust in the system’s proper and safe operation; organizational integrity – faith that the organization using AI is doing so responsibly; and relational authenticity – faith that automated communication still conveys genuine respect and intent. Together, these factors generate “synthetic trust”: a form of trust grounded in human values but enabled by technology, and therefore fragile – one mistake, one hallucination, or one opaque output can quickly erode both the AI system’s and the organization’s credibility.

Preserving trust in hybrid ecosystems, where AI systems generate recommendations, classifications, or communications, means guaranteeing stakeholders “the right to know why”: to find out that AI was involved, how it processed the relevant information, what data or parameters affected its output, and who is responsible for checking or approving that output. To operationalize this, organizations need to establish an AI stewardship, i.e., a disciplined way of employing intelligent technologies with care, prudence, and a sense of long-term responsibility, resting on three bases:

- Competence. Ensuring that everyone who uses AI systems knows what they can and cannot do – since misuse is often, simply, a matter of not knowing how to use a tool correctly.
- Integrity. Ensuring that every part of AI deployment and communication is transparent and accountable.

- Empathy. Always keeping in mind that there is a human story behind every dataset, and that every automated message reaches a person with hopes, fears, and feelings.

We cannot program trust, but we can plan for it, by adopting governance frameworks that keep people first and human oversight over automation. In the AI era, governance is multifaceted and essential, and it must address three interconnected domains:

- Ethical Governance. Setting fairness, non-discrimination, privacy, and human dignity as basic rules for AI behavior, supported by institutional control such as AI Ethics Committees or multidisciplinary review boards for sensitive uses.
- Operational Governance. Embedding AI use within regular project management processes, with the right checks and balances: AI suggestions must be reviewed by a human, and AI-related issues – bias, lack of transparency, misuse, security – must be part of the overall project risk register.
- Regulatory Governance. Ensuring compliance with evolving regulations on AI, data protection, and digital accountability, including GDPR, industry-specific rules, and international ethical guidelines.

A mature AI governance framework can be translated into five operational objectives, each with its own mechanisms:

- Transparency – ensuring stakeholders know when and how AI is used, through output labeling, documentation, and communication disclosure.
- Accountability – maintaining clear human ownership of AI outcomes, through role assignment, audit trails, and decision logs.
- Fairness – preventing discrimination and bias in AI behavior, through data auditing, ethical review, and diverse oversight teams.
- Trustworthiness – fostering reliable and respectful AI interactions, through validation testing, ethical charters, and feedback channels.
- Continuous learning – adapting governance to technological evolution, through periodic reviews, AI ethics training, and stakeholder feedback loops.

These mechanisms do not impede innovation; on the contrary, they promote it, since in the age of AI, legitimacy – not control – is what stakeholders want, and people and systems only operate well together when they believe that the methods that link them are fair. AI stakeholders, by their nature, cannot guarantee such fairness; only humans can. Governance is therefore not simply a set of constraints; it is the architecture of moral alignment that allows artificial and human agents to collaborate without violating the values that make projects meaningful and successful.

The Human Difference

Ethical decision-making itself becomes a two-step process in AI-augmented contexts. In the first step, AI uses available data and optimization logic to generate options; in the second step, people apply ethical reasoning, contextual understanding, and empathy to select the best one. This two-loop structure supports the principle of augmented judgment: AI grants individuals more analytical power, but people still hold moral judgment. No matter how sophisticated the algorithm, the final decision remains a human responsibility, and it should be based on values, not merely on probabilities. Consequently, training programs for project professionals need to evolve to include not only data and AI literacy, but also AI ethical literacy: the ability to question the fairness, transparency, explainability, and social effects of algorithmic technologies.

None of this diminishes the role of people; quite the opposite. Adding AI systems to the stakeholder ecosystem paradoxically reinforces the importance of humans rather than reducing it: having artificial actors in the network does not make responsibility less clear, it makes it clearer, because only humans in the network can think about morality, be responsible, and empathize. As communication, analysis, and even creativity become more automated, keeping a unique human touch becomes both a moral and a strategic necessity. This “human signature” is not recognized for its technological sophistication, but for its authenticity – the ability to act with care, humility, and meaning. In a world of increasingly capable synthetic voices, the real one stands out: being authentic becomes a competitive advantage, as much as a governance principle.

The next Articles of this series will explore how this new perspective operates in practice: how AI can support the identification and analysis of stakeholders through data mapping and sentiment interpretation; how it can improve engagement and communication while preserving authenticity; and how human-centered skills, governance, and ethics can ensure that the integration of AI stakeholders remains transparent, fair, and trustworthy. Ultimately, the opposite of artificial intelligence is not resistance, but human excellence – the capacity to combine rational thinking with moral sensitivity, to run systems while serving people, and to make sure that the search for efficiency never becomes a substitute for empathy.

The contents of this Article are extracted and synthesized from Chapter 15 of my book “The Stakeholder Perspective: Projects, People and Performance in the Generative AI Era” (2nd Edition), Taylor & Francis, CRC Press, Boca Raton (FL), U.S.A., June 2026, where Readers, if they wish, can find further insights and information.

REFERENCES

Bostrom N., 2014, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press.

Bryson J., 2019, *The Artificial Intelligence of the Ethics of Artificial Intelligence: An Introductory Overview for Law and Regulation*, The Hertie School of Governance.

Floridi L., 2019, *The Logic of Information: A Theory of Philosophy as Conceptual Design*, Oxford University Press.

Floridi L., 2023, *The Ethics of Artificial Intelligence: Principles, Challenges, and Opportunities*, Oxford University Press.

Pirozzi M., 2017, [The Stakeholder Perspective](#), Featured Paper, PM World Journal Vol. VI, Issue VI – June 2017.

Pirozzi M., 2019, *The Stakeholder Perspective: Relationship Management to Increase Value and Success Rates of Projects*, CRC Press, Taylor & Francis Group.

Pirozzi, M., 2024, [The Stakeholder Perspective in the Generative Artificial Intelligence Scenario and the AI-Stakeholders](#), PM World Journal, Vol. XIII, Issue VIII – August 2024.

Pirozzi, M., 2025, [AI-stakeholder: come gestire l'influenza nelle strategie aziendali.](#), Agenda Digitale.

Pirozzi M., 2026, *The Stakeholder Perspective: Projects, People and Performance in the Generative AI Era*, 2nd Edition, CRC Press, Taylor & Francis Group, Boca Raton (FL), U.S.A. ISBN 978-1-041-14837-1.

About the Author



Massimo Pirozzi

Rome, Italy



Massimo Pirozzi, MSc cum laude, Electronic Engineering, University of Rome “La Sapienza”, Project, Program, and Portfolio Manager, Lecturer, Educator, Generative AI Leader, and Specialist. He is a Member of the Scientific Committee and an Accredited Master Teacher of the Istituto Italiano di Project Management (Italian Institute of Project Management). He is certified as a Professional Project Manager, as a Generative AI Leader, as an AI Professional, as an Information Security Management Systems Lead Auditor, and as an International Mediator. He is specialized, among other specializations, in Generative Artificial Intelligence for Project and Program Managers, in Generative Artificial Intelligence for Educators, in Responsible Generative Artificial Intelligence, in Prompt Engineering for Project Managers and for Educators, and in Agentic AI. In general, he has more than 70 Credentials in Project Management, Artificial Intelligence, Emotional Intelligence, Creative and Critical Thinking, Learning Design and Education released by primary US, British, Italian, and Singaporean Universities, US Companies, US and Italian Professional Associations.

Massimo is a Researcher, a Lecturer, and an Author about Stakeholder Management, Relationship Management, Complex Projects Management, and Generative AI, and he authored or coauthored more than 50 works in 5 different Countries. In particular, he is the Author of the Book “[*The Stakeholder Perspective: Relationship Management to Enhance Project Value and Success*](#)”, CRC Press, Taylor & Francis Group, Boca Raton (FL), U.S.A., October 2019, the second Edition of which “[*The Stakeholder Perspective: Projects, People and Performance in the Generative AI Era*](#)”—including a consistent additional part dedicated to AI—has been published in June 2026. Due to the acknowledgement of his comments on stakeholder-related issues contained in the Exposure Draft of The Standard for Project Management—7th and 8th Editions, he has been recognized as one of the contributors and reviewers of the PMBOK® Guide—7th and 8th Editions, and he received the *Certificate of Appreciation for Excellence for his*

volunteer contributions to the Project Management Institute and the project management profession in 2020. He also received several International Awards.

Massimo Pirozzi has wide experience in managing large and complex projects, programs, and portfolios in national and international contexts, and in managing business relations with public and private organizations, including multinational companies, small and medium-sized enterprises, research institutes, and non-profit organizations. He worked successfully in several sectors, including Defense, Security, Health, Education, Engineering, Logistics, Cultural Heritage, Transport, Gaming, Services to Citizens, Consulting, and Web. He was also, for many years, a Top Manager in ICT Industry, and an Adjunct Professor in Organizational Psychology. He is registered as an Expert of both the European Commission and Italian Public Administrations.

Massimo Pirozzi is an Accomplished Author, a Member of the Executive Team and an International Editorial Advisor of *PM World Journal*, and can be contacted at max.pirozzi@gmail.com.

www.linkedin.com/in/massimo-pirozzi-the-stakeholder-perspective