Special Issue — R&D Management Journal

The Non-Human Innovator: Agentic AI, Physical AI, and the Transformation of R&D Management

Human–AI Delegation, Autonomous Agency, Organisational Capabilities, and AI Inventorship

Introduction & Rationale

Who is the innovator when the innovator is not human? This question — at once provocative and analytically precise — defines the organising challenge of this Special Issue. Open innovation, as Chesbrough (2003) defined it, assumes purposive human actors managing knowledge flows across organisational boundaries. Two decades of research in this tradition, much of it published in *R&D Management*, has examined how firms balance openness and appropriability (Dahlander & Gann, 2010), how knowledge breadth shapes innovation performance (Laursen & Salter, 2006), how open-source models create revealing incentives distinct from proprietary innovation (Von Hippel & Von Krogh, 2006), and how transitions from closed to open innovation require organisational restructuring (Chiaroni, Chiesa, & Frattini, 2010). Throughout this literature, the identity of the innovator — individual, organisations, or individuals, organisations, or consortia — has been taken as given: a human agent or collective of human agents.

Artificial intelligence (AI) now disrupts this foundational assumption. AI systems — including large language models(LLMs), autonomous agents, and physical AI technologies — are emerging not merely as tools for processing external knowledge but as active participants in the innovation process itself, shaping core R&D activities such as idea generation, experimentation, and development (Steinhauser & Heid, 2026). Recent work has advanced our understanding of generative AI across creativity, innovation performance, and organisational processes (Mariani & Dwivedi, 2024), but these insights remain conceptually fragmented and largely tool-centric. What is missing is a coherent framework for understanding AI as an emerging class of autonomous and semi-autonomous agents that reshapes innovation systems, redefines creativity as a hybrid socio-technical process, and introduces governance challenges that existing theory cannot resolve.

The deepest challenge is not how AI *assists* human innovators. It is how much of the cognitive, creative, and evaluative work of innovation should be *delegated* to non-human agents — and what happens to R&D organisations, ecosystems, and intellectual property (IP) regimes when that delegation becomes substantial. Laursen and Salter's (2006) inverted-U relationship between external search breadth and innovation performance raised the question of whether firms can be "over-open." We now face an analogous question about optimal delegation to AI. And just as Katz and Allen (1982) documented how insularity erodes R&D performance, AI-intensive organisations face a novel converse problem: an over-reliance on AI-generated knowledge that erodes the internal expertise and critical judgment required to evaluate and integrate it, potentially generating forms of "artificial certainty" that undermine expert authority (Leonardi & Leavell, 2026).

A further frontier emerges with agentic and physical AI: systems that do not merely assist human innovators but autonomously plan, execute, and iterate multi-step innovation tasks across physical and digital environments (Huang et al., 2024). The convergence of agentic AI with robotics, digital twins, and cyber-physical systems defines a regime in which intelligent systems act upon and reconfigure the physical world, raising new governance questions that the original open innovation framework was not designed to address. Gassmann, Enkel, and Chesbrough (2010) identified nine research perspectives needed for a mature open innovation theory; we contend that a tenth is now urgently required: the governance of non-human agency in innovation as an integral dimension of open innovation.

This Special Issue is explicitly positioned around a question that existing work on AI-enabled knowledge search, collaboration management, and general AI governance has not yet answered: what happens when AI ceases to be a mere instrument and becomes an agent of innovation? This distinction — between AI-as-tool and AI-as-innovator — defines the conceptual boundary of our four research themes, which address the optimal human–AI division of innovative labour, the architecture of autonomously populated innovation ecosystems, the reimagination of absorptive capacity for an agentic era, and the fundamental challenges to IP posed by non-human inventorship.

Research Themes

We welcome conceptual, empirical, and methodological contributions addressing, but not limited to, the following four themes.

Theme 1. How open should innovation be to AI? Human–AI Delegation and the Optimal Division of Innovative Labour

Focus: Extending the open innovation question of "how open is open enough" to the human–AI frontier: how much cognitive, creative, and evaluative work should be delegated to AI, at what stages of the R&D process, and under what governance conditions?

Laursen and Salter (2006) demonstrated that excessive external search can undermine integration capacity. Transposed to the human–AI relationship, this raises an equally pressing question: is there an optimal level of delegation to AI beyond which performance declines? As AI increasingly participates in a wide range of R&D activities, from hypothesis generation and experimental design to knowledge synthesis, R&D organisations must make consequential choices about where human judgment remains indispensable. AI also introduces a structural inversion of the NIH syndrome: organisations may face a "not-evaluated-here" (NEH) problem in which AI-generated knowledge may be accepted uncritically, bypassing the internal expert judgment that makes external knowledge genuinely valuable. Dodgson, Gann, and Salter's (2006) analysis of how technology reshapes open innovation processes at Procter & Gamble prefigures the transformations now underway as AI reconfigures the division of innovative labour.

Potential research questions:

- Is there an inverted-U relationship between the degree of AI delegation in R&D and innovation performance, mirroring the over-search paradox?
- How do firms determine the optimal breadth and depth of AI involvement across different stages of the R&D process? Specifically, under what conditions does algorithmic management enhance or hinder innovation performance?
- What is the AI-era analogue of the NIH syndrome: how does uncritical acceptance of AI-generated knowledge erode internal expertise?
- How must absorptive capacity be reconceptualised when the primary external knowledge source is an AI system rather than a human partner?
- What governance architectures enable effective human–AI teaming while preserving accountability, creativity, and strategic judgment?

Keywords: *Human–AI collaboration, optimal delegation, R&D governance, absorptive capacity, algorithmic management.*

Theme 2. Agentic AI and Physical AI as Autonomous Actors in Innovation Ecosystems

Focus: How agentic AI systems and physical AI technologies constitute fundamentally new actors in open innovation ecosystems — not tools, but autonomous agents — reconfiguring inter-firm collaboration, knowledge flows, and ecosystem governance in ways the original inbound/outbound/coupled typology could not anticipate.

Enkel, Gassmann, and Chesbrough (2009) identified three core open innovation processes in *R&D Management*: inbound, outbound, and coupled — all developed with human actors as the units of agency. Agentic AI systems capable of autonomously planning, executing, and iterating multi-step tasks introduce a fourth dimension: automated knowledge flows generated, processed, and transmitted by non-human agents operating continuously across organisational boundaries. Physical AI extends this logic further, raising questions about how firms govern inter-agent collaboration, how digital twins enable distributed experimentation, and how the integration of AI into manufacturing and service systems reshapes industry boundaries (Urbinati et al., 2020; Cappa et al., 2021). Understanding what "inbound," "outbound," and "coupled" open innovation means when AI agents are active participants is a fundamental question for the next generation of open innovation research.

Potential research questions:

- How do agentic AI systems alter the architecture of knowledge flows in open innovation ecosystems, and what new governance mechanisms are required?
- How does the participation of autonomous agents change the logic of platform-based open innovation, including roles, incentives, and boundary conditions?
- How do digital twins and AI-enabled simulation reshape the scope and speed of distributed experimentation across organisational boundaries?
- What new forms of inter-organisational trust, contracting, and coordination are needed when AI agents act as innovation partners?

Keywords: *Agentic AI, Physical AI, open innovation ecosystems, digital twins, platform governance, inter-organisational collaboration.*

Theme 3. Absorptive Capacity Reimagined: Organisational and Individual Capabilities for AI-Era Innovation

Focus: How the capability requirements for open innovation are being transformed by AI, including the redefinition of absorptive capacity, dynamic capabilities, and individual innovation competencies in human–AI environments.

Absorptive capacity — the ability to recognise, assimilate, and apply external knowledge for commercial ends (Cohen & Levinthal, 1990) — has been central to open innovation theory as an explanation of why firms differ in their ability to exploit external knowledge sources. In the era of AI, this concept requires fundamental reconceptualisation: the critical challenge is no longer whether firms can integrate human-generated external knowledge but whether they can effectively evaluate, critically interrogate, and recombine AI-generated knowledge — a qualitatively different cognitive and organisational challenge. Firms that developed strong absorptive capacity for human-generated scientific knowledge may lack the capabilities to critically assess large language model outputs, and may be systematically disadvantaged by overconfidence in AI-generated insights. Chiaroni, Chiesa, and Frattini's (2010) analysis of the transition from closed to open innovation provides a useful template for examining the analogous transformation now underway as firms open their innovation processes to AI participation (Haefner et al., 2021). These challenges also extend to dynamic capabilities –the "ability to integrate, build, and reconfigure internal and external competences to address rapidly changing environments" (Teece et al., 1997).

Potential research questions:

- How must absorptive capacity (or dynamic capabilities) be reconceptualised when the primary external knowledge source is an AI system rather than a human partner?
- What organisational routines and managerial processes enable firms to transform AI-generated knowledge into innovation value?
- What individual competencies — technical, cognitive, and relational — distinguish high-performing innovators in human–AI contexts?
- What learning mechanisms allow firms to continuously upgrade AI-related innovation capabilities over time, particularly in relation to exploration and exploitation?

Keywords: *Absorptive capacity, dynamic capabilities, AI literacy, organisational learning, exploration–exploitation, innovation competencies.*

Theme 4. AI Inventorship, Algorithmic Assets, and the Open-Source AI Dilemma: Who Owns What AI Creates?

Focus: How AI systems, as autonomous and semi-autonomous agents, create entirely new questions of inventorship, asset ownership, and strategic disclosure that existing IP frameworks were not designed to address — questions that are categorically distinct from the protection of human-generated innovations and from cybersecurity or appropriability mechanisms addressed in adjacent research.

The emergence of AI as a generative and autonomous agent of innovation forces a fundamental reexamination of who — or what — constitutes the legal and economic subject of IP. When agentic systems or autonomous laboratory robots generate patentable outputs with minimal human prompting, the attribution of inventorship becomes legally contested and theoretically unresolved. West and Gallagher's (2006) open-source paradox — developed for human developers deciding when to reveal versus protect software innovations — now applies with new intensity to foundation model ecosystems, where firms must decide whether to release model weights, training data, or fine-tuned capabilities into the commons. Von Hippel and Von Krogh's (2006) private-collective model, which theorised when free revealing is individually rational, requires extension to the non-human case. Trained model weights constitute a new category of strategic asset whose competitive value may dwarf that of any individual patent in a firm's portfolio, yet no existing IP instrument cleanly covers them. This theme explicitly excludes cybersecurity, trade secrecy enforcement, and appropriability mechanisms for human-generated innovations — focusing instead on the prior question: how should firms, legal systems, and innovation ecosystems reconceptualise IP when the generative agent is non-human?

Potential research questions:

- How should inventorship and IP ownership be attributed when agentic AI systems autonomously generate patentable outputs, and what theoretical frameworks from open innovation research best capture this challenge?
- What strategic logic governs firms' decisions to release model weights, training data, or fine-tuned AI capabilities into open-source commons?

- How do trained model weights, fine-tuning data, and emergent AI capabilities constitute a new category of strategic asset, and how do firms govern access to and monetisation of these assets?
- Under what regulatory and institutional conditions does the open-source AI movement accelerate versus impede innovation diffusion?

Keywords: *AI inventorship, model ownership, open-source AI dilemma, foundation model governance, algorithmic assets, non-human agency in IP.*

Submission Information

This Special Issue is partially associated with the R&D Management Workshop 2026 (Seoul). Participation in the workshop is not a prerequisite for submission; however, a paper development workshop (PDW) session will be held, offering authors presenting at the workshop the opportunity to discuss their drafts and enhance manuscript quality. We welcome conceptual, empirical, and methodological contributions from management, organisation science, and innovation studies.

Guest Editors

- Prof. Alberto Di Minin (Sant'Anna School of Advanced Studies)
- Prof. Joonmo Ahn (Korea University)
- Prof. Sungjoo Lee (Seoul National University)
- Prof. Ahreum Hong (Kyung Hee University)

Key Dates

- 29~30th May 2026: PDW[1] at R&D Management Workshop, Seoul, Korea
- 2~4th July 2026: PDW at KOSIME summer conference, Jeju, Korea
- 30th October 2026: Special Issue Submission Open
- 28th February 2027 Deadline for SI Submission
- 1st March 2028: Publication of SI Articles

Review Process

All submissions will undergo a standard double-blind peer-review process in line with *R&D Management* journal guidelines.

References

Adner, R. (2017). Ecosystem as structure: An actionable construct for strategy. Journal of Management, 43(1), 39–58.

Cappa, F., Del Sette, F., Hayes, D., & Rosso, F. (2021). Leveraging big data and open data for open innovation. R&D Management, 51(5), 446–461.

Chesbrough, H. (2003). Open Innovation: The New Imperative for Creating and Profiting from Technology. Harvard Business School Press.

Chiaroni, D., Chiesa, V., & Frattini, F. (2010). Unravelling the process from closed to open innovation. R&D Management, 40(3), 222–245.

Cohen, W. M., & Levinthal, D. A. (1990). Absorptive capacity: A new perspective on learning and innovation. Administrative Science Quarterly, 35(1), 128–152.

[1] PDW(paper development workshop): There will be two optional PDWs at Seoul (May 2026) and Jeju (July 2026). The attendance of the PDW is not mandatory.

Dahlander, L., & Gann, D. M. (2010). How open is innovation? Research Policy, 39(6), 699–709.

Dodgson, M., Gann, D., & Salter, A. (2006). The role of technology in the shift towards open innovation: The case of Procter & Gamble. R&D Management, 36(3), 333–346.

Enkel, E., Gassmann, O., & Chesbrough, H. (2009). Open R&D and open innovation: Exploring the phenomenon. R&D Management, 39(4), 311–316.

Gassmann, O., Enkel, E., & Chesbrough, H. (2010). The future of open innovation. R&D Management, 40(3), 213–221.

Haefner, N., Wincent, J., Parida, V., & Gassmann, O. (2021). Artificial intelligence and innovation management. Technological Forecasting and Social Change, 162, 120392.

Huang, J., et al. (2024). Understanding the planning of LLM agents: A survey. arXiv preprint arXiv:2402.02716.

Jacobides, M. G., Cennamo, C., & Gawer, A. (2018). Towards a theory of ecosystems. Strategic Management Journal, 39(8), 2255–2276.

Katz, R., & Allen, T. J. (1982). Investigating the not-invented-here (NIH) syndrome. R&D Management, 12(1), 7–19.

Laursen, K., & Salter, A. (2006). Open for innovation: The role of openness in explaining innovation performance. Strategic Management Journal, 27(2), 131–150.

Laursen, K., & Salter, A. J. (2014). The paradox of openness: Appropriability, external search and collaboration. Research Policy, 43(5), 867–878.

Leonardi, P. M., & Leavell, V. (2026). Knowing enough to be dangerous: The problem of "artificial certainty" for expert authority when using AI for decision making and planning. Organization Science, 37(2).

Lichtenthaler, U., & Ernst, H. (2006). Attitudes to externally organising knowledge management tasks. R&D Management, 36(4), 367–386.

Mariani, M., & Dwivedi, Y. K. (2024). Generative AI in innovation management: A systematic review and future research agenda. Journal of Business Research, 172, 114447.

March, J. G. (1991). Exploration and exploitation in organizational learning. Organization Science, 2(1), 71–87.

Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management. Academy of Management Review, 46(1), 192–210.

Steinhauser, T., & Heid, B. (2026). Organizational readiness, AI literacy, and the new frontier of R&D: How generative AI shapes innovation capacity. R&D Management.

Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. Strategic Management Journal, 18(7), 509–533.

Teece, D. J. (1986). Profiting from technological innovation. Research Policy, 15(6), 285–305.

Teece, D. J. (2018). Profiting from innovation in the digital economy. Research Policy, 47(8), 1367–1387.

Urbinati, A., Chiaroni, D., Chiesa, V., & Frattini, F. (2020). Digital technologies and open innovation. R&D Management, 50(1), 136–160.

Von Hippel, E., & Von Krogh, G. (2006). Free revealing and the private-collective model for innovation incentives. R&D Management, 36(3), 295–306.

West, J., & Gallagher, S. (2006). Challenges of open innovation: The paradox of firm investment in open-source software. R&D Management, 36(3), 319–331.

Zahra, S. A., & George, G. (2002). Absorptive capacity: A review, reconceptualization, and extension. Academy of Management Review, 27(2), 185–203.