

## Scaling Ai Features in Large Organizations: A Product Management Perspective

Obianuju Gift Nwashili\*

Independent researcher

| Corresponding Author:  | Obianuju Gift Nwashili<br>Independent researcher                                  | <b>Abstract:</b> Scaling AI capabilities from a promising Proof-of-Concept (POC) to a widely adopted, production-ready product has become one of the most important and complex organizational challenges of our time. Recent studies have indicated a failure rate of over 80% for AI projects not making it past the pilot phase and into scaled production, resulting in vast amounts of talent and resources being consumed with no value delivered to the organization. This comprehensive review will serve as an expansive guide to help product managers develop a pragmatic, tactical approach to the “scaling AI” problem. We believe that scaling AI to production is first and foremost a product-led orchestration problem. AI scaling is a multi-faceted problem that must be solved in parallel with respect to “bleeding edge” technology and proven business value, operational maturity and cross-functional alignment. The framework shared here describes a four-phase lifecycle (Strategic Pilot, Operational Crucible, Cross-Company Scaling, Monetization) where the product manager needs to “own the whole stack” of the execution in order to methodically de-risk scaling. The product manager is the chief integrator and orchestrator of technical feasibility, human-centric design, business strategy and operational pragmatism. The goal is to productize AI to transform it from an interesting science experiment to a sustainable core differentiator and engine of profit for the company. |
|------------------------|---|---|
| <b>Article History</b> | Received: 22 / 10 / 2025<br>Accepted: 02 / 12 / 2025<br>Published: 10 / 12 / 2025 | <b>Keywords:</b> <i>AI Scaling; Product Management; Human-AI Collaboration; Enterprise AI Adoption; AI Productization.</i>  |

**How to Cite in APA format:** Nwashili, O. G. (2025). Scaling Ai Features in Large Organizations: A Product Management Perspective. *IRASS Journal of Economics and Business Management*. 2(12), 23-30.

### Introduction

Organizations that achieve enterprise agility and adaptability to adopt AI and ML technologies build organizational strength while driving creative innovation and maintaining strategic advantages (Abonamah & Abdelhamid, 2024). In fact, to what extent an organization can do this becomes a major point of differentiation. In order to win, however, this requires a product philosophy shift and reorientation. It's imperative for organizations to deeply appreciate and internalize the product side of AI in two distinct but interconnected facets: AI not just as a potential redefinition of product-market success but also as a radical revision of how products are built and evolved (Witkowski & Wodecki, 2025). With many AI initiatives getting stuck in what has been wittily referred to as “pilot purgatory” or the black box between proof-of-concept and enterprise-scale production-grade success, most large-scale companies, from startups to enterprises are floundering to convert their AI/ML efforts to scaled production realities.


Scaling AI initiatives face a range of barriers that create a sizable distance between ambitious experimentation and proof-of-concept (POCs) and the realities of real-world production-level deployment and value creation (Rumalla & Mujawar, 2025). This includes the high financial cost of large-scale implementation, which in some cases involves steep technology integration costs to new software tools, data pipelines, and specialized hardware. Additionally, it also includes new skill and talent acquisition (Kasireddy, 2025). However, these challenges often pale in comparison to a set of human-centric constraints such as workforce resistance that stems from real or perceived threats to job security, AI-related skill anxiety, or a general lack of cultural acceptance within existing workflows and legacy

processes. Increasingly, studies point towards end-user acceptance, contextual utility, and compatibility with existing workflows often emerging as more decisive in adoption success than merely outperforming raw technical metrics (Jacob, 2025). The recent hype cycle around Generative AI has led to a spike in trials but it has also led to a corresponding increase in unreliable execution. One line of recent research shows that 25-50% of AI research and development projects simply fail or fizzle, often due to failure to create outcomes that can be reused or reproduced in operationally scaled ways (Herremans, 2021; Rumalla & Mujawar, 2025).

In a concerted effort to offer a clear roadmap to successful at-scale adoption, this guide takes a direct approach to solve the more complex socio-technical conundrums of AI scaling by centering AI scaling as a core product management concern and challenge. The guide, in a deliberate countermove to an overly technical implementation view, approaches AI adoption as an approach to strategic product innovation and repositioning rather than a product optimization or efficiency play (Climent et al., 2024; Nikolić & Bjelica, 2025). This approach enables AI adoption to create immediate quantifiable value, with the guide's core framework uniquely created to solve a common gap product managers face — a mismatch between executive level aspirations and day-to-day operational requirements. The guide offers a roadmap and toolkit to reduce this often very costly and risky misalignment, with the promise of creating effective AI-backed value (Nikolić & Bjelica, 2025).

The product manager role itself is also in the middle of a sea change, with Parikh (2025) labeling this the “age of agentic AI” where the product manager can no longer view their role as

one of linearity, producing features but will need to expand their purview and focus to AI-first strategy and the creation and governance of expansive product socio-technical systems (STES). It also points to a concomitant expansion of the very practice of product management from an old “predict-and-plan” mindset and how product management is usually taught and applied to a more “adaptive hypothesis” mindset, to reflect the epistemic uncertainty of such powerful generative and “next-gen” AI tools (Olsson & Bosch, 2024). This guide is part roadmap, part curated toolkit for the product manager, to structure their approach in the following specific ways that then powerfully open up the possibility of actualizing AI’s radical potential while also methodically and predictably taking a feature from proof-of-concept to tangible, scaled value.



**Figure 1.** AI adoption outcomes illustrating the steep attrition from proof-of-concept to scaled production deployment.

## Literature Review

### The Transformative Impact of AI on Product Management and Innovation

Artificial intelligence is reshaping product management as a discipline, in particular the New Product Development process, like no other technology. The value-add of AI has thus far been most tangible in terms of amplifying and accelerating existing product management processes: applications in rapid prototyping, in automated analysis of natural language customer feedback, as well as in AI-assisted code generation have all been shown to reduce time-to-market and cost of development without sacrificing (and often with an increase in) product-market fit (Ateeq et al., 2025; Parikh, 2023). The entire process of innovation has become more fluid, iterative, and data-driven, with an increasing need to transform innovation management accordingly (Füller et al., 2024). In particular, the rise of Generative AI is set to be a game changer.

Generative AI, enabled by LLMs, diffusion models and other emerging models is a formative innovation in itself. Its generative AI models are being deployed as part of the innovation process, with unique power to condense the typically fuzzy front end of the innovation process—i.e., the ideation process and compress the time and cost of experiments and concept validation at extreme levels (Corvello, 2024; Mariani & Dwivedi, 2024; Roberts & Candi, 2024). Inevitably, this rapid development and deployment of AI has also created major ethical, practical and strategic challenges on an urgent timescale, and many of these are directly related to product managers’ work. Chief among these

responsibilities is the (non-delegable) responsibility of product managers to set ethical guardrails and other governance to ensure responsible AI usage. There is a growing tension between the need to innovate rapidly and capture value, and to build AI in a way that is safe, fair, and socially responsible (Smith et al., 2025).

### The Imperative of Human-AI Co-Creation

An emerging body of literature is making a case for a human-centric, co-creative approach as the most productive framework for deploying Generative AI. Methodologies such as the AI-augmented double diamond framework provide a blueprint for integrating AI as a ‘fourth partner’ in human-centred design and strategy processes. They outline the value of AI as an enhancer of human capabilities for a range of NPD tasks, from automating text summaries and systematising concept evaluation, to interpreting qualitative feedback and generating divergent concepts (Bouschery et al., 2023).

The most productive and sustainable use cases in practice involve Generative AI less as a self-sufficient problem-solver and more as an additional ‘member’ of a human-AI co-creation loop. In this setup, Generative AI serves more as an amplifier and collaborator to human creativity and critical thinking than as its substitute. Evidence of the effectiveness of this approach is emerging from empirical studies. For instance, it has been reported that human-AI co-creation teams significantly outperformed their human-only counterparts in terms of both the quantity and novelty of ideas generated, particularly when diverse AI interaction and search strategies were utilised (Boussouix et al., 2024; Liu, 2025). This line of evidence is telling: it shows that the real value lies not in AI or humans working separately but in how they work in concert. The challenge for firms and businesses, then, is in how to structure and create the conditions for this to happen, i.e. to effectively orchestrate the alignment between human intuition, domain expertise, and ethical judgment on the one hand and the computational power and pattern recognition of AI tools on the other (Füller et al., 2024).

### The Evolving Research Landscape: From Efficiency to Transformation

To position my research, we reviewed three recent literature reviews about the broader theme of innovation management and AI. First, Mariani and Dwivedi (2024) have conducted a scoping review of the literature on innovation and AI in management. In their findings, they note that the research is “mature” (p.6) and has shifted from AI use in a narrow efficiency or automation sense to a broader transformational perspective of how products are designed, for whom, and with what value capture.

In particular, generative AI is often mentioned as a key technology to apply in this new way of managing innovation, which Corvello (2024) also mentions, even using the term “meta-innovation” (the ability to innovate new innovation approaches) in his recent review of innovation using large language models.

In addition to these shifts in the innovation topic, they note that managers will need “ambidextrous” (Holmström & Carroll, 2024) capabilities in both exploitative (AI for continuous improvement) and exploratory (AI for radical or business model innovation) forms, as the emerging AI tools are now more accessible (for non-specialists) through no-code/low-code platforms (Sherson, 2024). This produces the actionable problem

for practitioners that product managers will need a robust and practical framework for AI integration into innovation workflows to foster effective human-AI collaboration for the best results.

## Methodology

This framework is derived from a comprehensive, multi-pronged, and methodologically rigorous inquiry that strikes a balance between academic rigor and real-world applicability. At its core is a systematic literature review that scoured the theoretical landscape, charted dominant trends, and identified key gaps in the convergence of AI scaling, product management, and organizational change literature (Hemraj, 2025; Wang et al., 2023). This review included both academic journals and conference proceedings, as well as key industry whitepapers and reports.

To ground theory in the complex realities of organizational life, we also incorporate rich qualitative data from multiple case studies and semi-structured interviews with experienced product managers, AI platform leaders, and digital transformation executives across different industries. This qualitative component allows for the identification of unmet needs, tacit challenges, and emerging (yet often uncodified) best practices, which are less visible in survey data or standardized frameworks (Cooper & Brem, 2024).

Qualitative findings are then tested and generalized with quantitative survey data from a large sample of product development and technology professionals. This quantitative analysis seeks to establish statistical relationships between specific AI adoption patterns, team compositions, governance models, and key performance metrics such as time-to-market, feature adoption rates, and ROI (Ogundipe et al., 2024).

Finally, we integrate cutting-edge experimental research on human-AI collaboration. This includes controlled studies that assess the outcomes, processes, and team dynamics of human-AI co-creation teams versus human-only teams. This experimental data offers valuable insights into the tangible effects of AI on cognitive processes, creative ideation, solution quality, and implementation success across different phases of the NPD lifecycle (Boussouix et al., 2024; Hou et al., 2025; Lee & Maruping, 2024).

By converging insights from these diverse research strands, we ensure that the resulting framework not only achieves depth and breadth but is also deeply rooted in both the theoretical and practical realities of scaling AI technologies in complex, matrixed, large organizations.




Figure 2. Four-phase product-led framework for scaling AI from strategic pilot to monetization and sustained value.

### PHASE 1: The Strategic Pilot – Laying the Unshakeable Foundation

The first principle is to start with a strategic pilot: small-scale production deployments have always been the starting point for scaling but have often been misdirected. A successful scaling effort requires a pilot that is product-focused, hypothesis-driven and planned from the beginning with a future of scale in mind. The first key step for the product manager is therefore to choose a use case with a clear, direct line to measurable business outcomes. Sweeping statements about “better insights” must be converted to concrete goals and measures: “10% reduction in customer service handling time” or “5-point increase in checkout conversion”, paired with one primary North Star metric on which success will be evaluated throughout the scale journey. Most importantly: this pilot should not be a “sandbox” effort. It must take place in the actual target deployment environment using real production data (not proxy or synthetic datasets), with real end-users and directly integrated with live backend systems. Only in this “real world” pressure test will the secret dependencies, data pipeline or integration complications that make-or-break scaling efforts later come to light.

### PHASE 2: The Operational Crucible – The Discipline of Productization

This stage is where the shift from “model-centric” PoC to a “product-ready” capability needs to occur. It is the stage where

most AI projects derail in their quest to the last mile of operationalization.

- **Infrastructure & Architecture:** The product manager must work with data science and engineering teams to elevate interesting capabilities from one-off scripts or notebooks to a scalable, reliable, and auditable MLOps pipeline. This means an industrialized workflow covering the entire lifecycle: from ingestion and validation of robust and versioned training data to model training and serving, to performance and drift monitoring, and automated CI/CD and seamless integration of AI/ML updates
- **Responsible AI & Governance:** Responsible AI must be productized as part of the product manager’s feature specifications. This includes building guardrails for ethical AI use (active fairness for bias mitigation, explainability, data stewardship, and security) directly into the AI feature from the ground up. Model cards, detailed audit trails, and even compliance checkpoints should be productized as part of the minimum viable product (MVP) deliverables to avoid late-stage, organizationally-damaging rework and necessary adjustments to align with outside regulations
- **User Experience & Change Management:** Designing the UX for an AI feature is also about managing the “trust calibration” users will have with that feature. Users must be given appropriate transparency into confidence

and decision-making by the AI and, more importantly, clear and frictionless human oversight and override controls for the AI feature. Human resistance to “black box” automation and concerns of skills obsolescence are

to be expected and should be met by proactive user communication and supporting HR and change management efforts to help shepherd the organization to the other side.

**Table 1. Core deliverables aligned to each phase of the AI scaling lifecycle.**

| Phase                    | Key Product Manager Deliverables                  | Success Indicators                    |
|--------------------------|---|---------------------------------------|
| 1. Strategic Pilot       | Real-data pilot, North Star Metric, UX validation | Early adoption, feasibility confirmed |
| 2. Operational Crucible  | MLOps pipeline, Model cards, HITL controls        | Reliability, governance, trust        |
| 3. Cross-Company Scaling | Shared platform, reusable components, CoE         | Reuse rate, faster next case          |
| 4. Monetization          | Business case, ROI tracking, A/B learning loops   | Revenue impact or cost savings        |

### PHASE 3: Cross-Company Scaling & Integration – The Organizational Lever

Scaling within a large enterprise organization is as much an “organizational design” challenge as a technical one.

- **The Platform Mindset:** In the name of efficiency and avoiding “AI sprawl,” the product manager has to be a product “platform evangelist,” lobbying for and contributing to common shared AI platforms, reusable component toolkits, standard libraries, and codified open APIs. Central to this perspective is a commitment to standardization and code reuse, which will also decrease product fragmentation and speed up development velocity on follow-on use cases.
- **Stakeholder Alignment & Funding:** Another essential enabler of success is an ongoing commitment to a high state of investment. The product manager must be able to craft a fact-based, data-driven business case that will pull the work forward out of its typical (inevitable) initial R&D/proof-of-concept “silo” and into the product-line or business-unit mainstream. This business case must make an unassailable ROI, strategic alignment, and value-capture case to stakeholders, including finance, executive management, and business-unit leaders.
- **Talent & Center of Excellence:** A good scaling model works in a hub and spoke pattern. A product manager needs to advocate for and work with a centralized (AI/ML) Center of Excellence (CoE), which can play an anchor role for the organization by setting common standards, providing deep subject-matter support, and also serving as an overall “community of practice.” At the same time, a product manager must “empower the edge,” to “grow their own” local, function-specific, and/or product-line-specific AI talent and teams at the business-unit level and project-manage them to build iteratively and constructively on these common starting points, not away from them.

### PHASE 4: Monetization & Sustained Value – The Business Endgame

The final phase marks the transition from enabling capability to value capture and growth.

- **Value Capture Models:** It is not enough to state “AI adds value”. The product manager must operationalize the value capture model by precisely articulating how it benefits the bottom line. Is it driving new revenue (through premium tiers, new markets, or price optimization) or protecting and amplifying profit (through cost avoidance, accelerated time-to-value, risk reduction, or operational efficiency)? Quantifying this value in business terms is a prerequisite for long-term success.
- **Continuous Learning & Evolution:** The deployed AI model is not a fixed artifact but a “moving target” that decays if not nurtured. The product manager must put in place regimens of continuous learning: A/B testing to optimize performance, feedback loops to capture user corrections, and diligent monitoring for data and concept drift. The AI product must be architected to adapt and evolve over time in line with shifting business objectives and user behaviors.
- **Managing the Portfolio:** As the organization builds multiple AI capabilities, the product manager must take a portfolio view. This includes deciding in what order to scale which features next, making portfolio investment decisions informed by a trade-off analysis of expected business value, technical and operational risk, organizational readiness, and strategic fit with the broader product ecosystem.


## Discussion

### Navigating the Paradoxes of Human-AI Partnership

Generative AI at scale magnifies all of the principles above in a differentially positive and a differentially negative way. On the one hand, it puts in the hands of people organizational forms of matter and mind with productivity levers unlike anything else in human history (Zysman, 2024). On the other, it gives people and organizations sources of error and failure with pathways to


propagate in and through the organization (Zysman & Nitzberg, 2024). If harnessed well, these new productivity levers are nothing

short of transformative. But there are equally transformational and non-obvious new sources of failure associated with them.



**Figure 3.** Desired trajectory where human expertise and AI capability reinforce each other rather than substitute one another

The fundamental reason is that even the most powerful LLMs today are (1) brittle and (2) inaccurate in meaningful ways, because of the pathologies of the LLMs that we know: hallucinations, systematic errors in reasoning, and all the other biases (Zysman & Nitzberg, 2024). The result is that any high-stakes operationalization of LLMs requires robust human-in-the-loop (HITL) validation architectures not as a “training wheels” interim stage, but as a permanent feature of the solution, both to secure the accuracy of output and to maintain and build user trust and to operationalize risk controls (Zhong, 2024). This is what Chin et al. (2024) call the “inferential trilemma” users find themselves in, of having to assess whether each output is a true AI “magic” moment to accept and act upon; a “hallucination” to be corrected; or a proposal “off the table” because of commercial or ethical concerns.



**Figure 4.** The inferential trilemma: Productivity gains vs. new failure pathways in generative AI systems.

The problem is that this changes the role of the human not merely from content-generator to content-evaluator but changes the kind of thinking required to do the role and hence the profile of human users needed to do it. There are dark sides to this re-skilling too. If this validation phase is not done with design, re-skilling human evaluators of LLM outputs can rapidly turn into deskilling, where users' cognitive abilities are replaced by the suggestions of the tool, leading to a rapid atrophy of the users' domain expertise

and critical faculties. Bastani et al. (2025) show this to be true in the legal context, and Sarkar et al. (2024) find similar evidence more broadly. In many cases, the key human cognitive struggle becomes not to generate answers, but to evaluate them (Drosos et al., 2025), even as it is often far more difficult, time-consuming, and technically unsupported.

This has a related well-understood dark side effect of cognitive automation bias, where the simple use of a computational tool for answers makes the user less able to accept them are not the right answer, when they clearly are not. This “bias” is the predictable psychological effect of users coming to expect or defer to the tool to have the right answer and not always being able to overcome the automation recommendation in the face of contradictory data or even their own better (latter-day) judgment.

Simkute et al. (2024a) call these various effects the “ironies of automation.” Automation has a long history of having these effects. Making highly automated assistants can make it even easier to do simple or routine tasks. But it also makes simple tasks incredibly easy to do and makes it much harder for any errors in judgment, often at a very deep, metacognitive level of “thinking about the thinking” that one is doing. If this metacognitive work is not done well by design, as Sarkar et al. (2024) find, highly automated assistants also have the potential to make it harder to do complex, uncertain, or high-variance tasks better or with greater domain expertise and instead to replace critical thinking skills of users with a kind of “magical thinking” around the tool.

Human-AI interfaces thus must be designed metacognitively as well as intuitively. Design can and must be used to drive critical thinking, to enhance not just simple work, but also metacognitive work. Promising lines of work in the recent work on this problem include the use of “provocations” to metacognitive work, by having the AI system generate divergent views of the same problem or solution, to surface likely uncertainties or sources of bias in the output or counter-questions to the user (Drosos et al., 2025; Singh et al., 2025). In effect, these are cognitive interrupt signals built into the solution to engage and

re-engage the user's metacognitive work and avoid their falling into simple or patterned behaviors around the solution. The design challenge of human–AI interfaces at scale is thus no longer simply one of reducing user effort, but one of calibrating user metacognitive engagement.


## Conclusion

### Extraheric AI for Scalable Product Management

The central thesis of this review culminates in a call for a fundamental philosophical shift in how we conceive of and design AI systems for scale. The key design objective should not be an AI solution that merely automates tasks but one that truly augments and extends human expertise, rather than potentially degrades or displaces it (Tankelevitch et al., 2024). To this end, we must evolve beyond AI tools that merely give answers and toward systems that empower and enhance thinking.

Our recommendation includes establishing a new design philosophy for extraheric AI systems which derives its name from the Latin term *extrahō* meaning "to draw forth" (Yatani et al., 2024). An extraheric AI system is designed not as a black box oracle, but rather as a true Socratic partner. Its core functionality shifts from one of simply "giving the right answer" to "drawing out" the human user's own critical thinking, judgment, and metacognition. It may take the form of an AI assistant that, when directly queried for an answer, instead offers carefully considered probing questions to help the user better scope and conceptualize the problem at hand, or that offers three plausible hypotheses each with their attendant limitations, or that deliberately challenges the user's opening assumptions with a few telling counterexamples. Its value lies in its ability to "draw forth" the human's higher-order reasoning.

This philosophy also directly addresses what Chen et al. (2025) call the "Assistance Dilemma": the counterintuitive reality that the more assistance an AI system provides, the less cognitive engagement it may actually elicit from the human, with the predictable downstream effect that the human may learn and develop less expertise as a result—precisely the opposite of the original intent of providing that assistance in the first place. The solution is a system of dynamic, elevated cognitive scaffolding (Riva, 2025). In such a system, the degree, nature, and even duration of the AI's scaffolding support is no longer static but is dynamically and tightly calibrated to the user's demonstrated competence, engagement, and the task's cognitive demands. Like training wheels on a bike, this scaffolding would be intentionally programmed to dissipate as user proficiency increases, automatically and actively shifting agency and cognitive load back to the human. This ensures the human remains the "expert in the loop" whose judgment and skills are being actively strengthened, not just left to atrophy from disuse.



**Figure 5.** Extraheric AI design model to strengthen user metacognition through dynamically adaptive cognitive scaffolding.

What this all means for the product manager leading the charge on AI scaling is nothing less than a recalibration of their north stars. It requires not only building product roadmaps and defining success metrics that place a premium on human skill-building and decision-making quality in addition to traditional efficiency measures but also partnering with UX researchers and designers to create interfaces that are as thought-provoking as they are easy to use. It requires also advocating for company-wide training programs to build not just AI technical literacy but true AI fluency as a critical component of the modern professional's judgment.

In summary, the successful scaling of AI at scale is the product management challenge of the decade. It requires a "whole stack" approach that fluently straddles technical architecture, ethical governance, org dynamics, business model innovation, and a core mastery of human cognition and collaboration. By following the phased approach and adopting the extraheric, human-amplifying design principles outlined in this handbook, product managers can turn the otherwise perilous, inefficient leap from POC to profit into a disciplined, de-risked process. In doing so, they will not simply be shipping scalable AI features but also will help build more adaptive, expert, and intelligently augmented organizations that can truly turn the massive potential of artificial intelligence into sustainable competitive advantage.

## References

1. Abonamah, A. A., & Abdelhamid, N. (2024). Managerial insights for AI/ML implementation: A playbook for successful organizational integration. *Discover Artificial Intelligence*, 4(1). <https://doi.org/10.1007/s44163-023-00100-5>
2. Ateeq, K., Masaeid, T. A., Selim, H., Oswal, N., Alkubaisy, A. A. A., Alami, R., & Ajdoobi, S. M. G. A. (2025). Harnessing AI for faster innovation: How AI concept generation impacts development timelines and market agility. *Journal of Posthumanism*, 5(2), 37. <https://doi.org/10.63332/joph.v5i2.390>
3. Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakci, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of*

*Sciences*, 122(26). <https://doi.org/10.1073/pnas.2422633122>

4. Bouschery, S. G., Blažević, V., & Piller, F. T. (2023). Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management*, 40(2), 139–153. <https://doi.org/10.1111/jpim.12656>
5. Boussioux, L., Lane, J. N., Zhang, M., Jaćimović, V., & Lakhani, K. R. (2024). The crowdless future? Generative AI and creative problem-solving. *Organization Science*, 35(5), 1589–1609. <https://doi.org/10.1287/orsc.2023.18430>
6. Chen, X., Ruan, K., Ju, K. P., Yap, N., & Wang, X. (2025). More AI assistance reduces cognitive engagement: Examining the AI assistance dilemma in AI-supported note-taking. \*Proceedings of the ACM on Human-Computer Interaction, 9\*(7), 1–24. <https://doi.org/10.1145/3757632>
7. Chin, T., Ghouri, M. W. A., Jin, J., & Deveci, M. (2024). AI technologies affording the orchestration of ecosystem-based business models: The moderating role of AI knowledge spillover. *Humanities and Social Sciences Communications*, 11(1). <https://doi.org/10.1057/s41599-024-03003-7>
8. Climent, R. C., Haftor, D., & Staniewski, M. W. (2024). AI-enabled business models for competitive advantage. *Journal of Innovation & Knowledge*, 9(3), 100532. <https://doi.org/10.1016/j.jik.2024.100532>
9. Cooper, R. G., & Brem, A. (2024). The adoption of AI in new product development. \*Research-Technology Management, 67\*(3), 44–52. <https://doi.org/10.1080/08956308.2024.2324241>
10. Corvello, V. (2024). Generative AI and the future of innovation management: A human-centered perspective and an agenda for future research. *Journal of Open Innovation: Technology, Market, and Complexity*, 10, 100456. <https://doi.org/10.1016/j.joitmc.2024.100456>
11. Drosos, I., Sarkar, A., Xiaotong, & Toronto, N. (2025). "It makes you think": Provocations help restore critical thinking to AI-assisted knowledge work. arXiv. <https://doi.org/10.48550/arXiv.2501.17247>
12. Füller, J., Tekić, Ž., & Hutter, K. (2024). Rethinking innovation management—How AI is changing the way we innovate. *The Journal of Applied Behavioral Science*, 60(4), 603–627. <https://doi.org/10.1177/00218863241287323>
13. Hemraj, S. (2025). AI and the future of creative development: Redefining digital media production. *AI and Ethics*. <https://doi.org/10.1007/s43681-025-00765-x>
14. Herremans, D. (2021). *AIStrom — A roadmap for developing a successful AI strategy*. arXiv. <https://doi.org/10.48550/arXiv.2107.06071>
15. Holmström, J., & Carroll, N. (2024). How organizations can innovate with generative AI. *Business Horizons*, 68(5), 559–570. <https://doi.org/10.1016/j.bushor.2024.02.010>
16. Hou, J., Wang, L., Wang, G., Wang, H. J., & Yang, S. (2025). The double-edged roles of generative AI in the creative process: Experiments on design work. *Information Systems Research*. <https://doi.org/10.1287/isre.2024.0937>
17. Jacob, G. (2025). From design to delivery: The strategic role of product managers in deploying AI solutions for patient-centered healthcare. *International Journal of Scientific Research in Science and Technology*, 12(5), 467–475. <https://doi.org/10.32628/ijsrst2512554>
18. Kasireddy, L. C. (2025). Overcoming adoption barriers: Strategies for scalable AI transformation in enterprises. *Journal of Informatics Education and Research*, 5(2). <https://doi.org/10.52783/jier.v5i2.2459>
19. Lee, H., & Maruping, L. M. (2024). *The rise of AI-enabled new product ideation: End of innovation crowdsourcing?* SSRN. <https://doi.org/10.2139/ssrn.5071267>
20. Liu, Z. (2025). *Human-AI co-creation: A framework for collaborative design in intelligent systems*. arXiv. <https://doi.org/10.48550/arXiv.2507.17774>
21. Mariani, M. M., & Dwivedi, Y. K. (2024). Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research*, 175, 114542. <https://doi.org/10.1016/j.jbusres.2024.114542>
22. Nikolić, M., & Bjelica, D. (2025). Digital product success under the microscope: When artificial intelligence in projects helps — and when it hurts. *PLoS ONE*, 20(8). <https://doi.org/10.1371/journal.pone.0331229>
23. Ogundipe, D. O., Babatunde, S. O., & Abaku, E. A. (2024). AI and product management: A theoretical overview from idea to market. *International Journal of Management & Entrepreneurship Research*, 6(3), 950–967. <https://doi.org/10.51594/ijmer.v6i3.965>
24. Olsson, H. H., & Bosch, J. (2024). Strategic digital product management in the age of AI. In *Lecture Notes in Business Information Processing* (Vol. 508, pp. 344–358). Springer. [https://doi.org/10.1007/978-3-031-53227-6\\_24](https://doi.org/10.1007/978-3-031-53227-6_24)
25. Orugboh, O. G., Omabuwa, O. G., & Taiwo, O. S. (2024). Predicting Neighborhood Gentrification and Resident Displacement Using Machine Learning on Real Estate, Business, and Social Datasets. *Journal of Social Sciences and Community Support*, 1(2), 53–70.
26. Orugboh, O. G., Omabuwa, O. G., & Taiwo, O. S. (2025). Predicting Intra-Urban Migration and Slum Formation in Developing Megacities Using Machine Learning and Satellite Imagery. *Journal of Social Sciences and Community Support*, 2(1), 69–90.
27. Parikh, N. A. (2023). Empowering business transformation: The positive impact and ethical considerations of generative AI in software product management — A systematic literature review. arXiv. <https://doi.org/10.48550/arXiv.2306.04605>
28. Parikh, N. A. (2025). *Agentic AI in product management: A co-evolutionary model*. arXiv. <https://doi.org/10.48550/arXiv.2507.01069>
29. Riva, G. (2025). *The architecture of cognitive amplification: Enhanced cognitive scaffolding as a resolution to the comfort-growth paradox in human-AI cognitive integration*. OSF Preprints. <https://doi.org/10.31219/osf.io/4wcdb>

30. Roberts, D., & Candi, M. (2024). Artificial intelligence and innovation management: Charting the evolving landscape. *Technovation*, 136, 103081. <https://doi.org/10.1016/j.technovation.2024.103081>

31. Rumalla, S., & Mujawar, J. S. (2025). Strategic decision-making framework for evaluating and selecting GenAI use cases. *International Journal of Science and Research Archive*, 16(1), 170–182. <https://doi.org/10.30574/ijrsa.2025.16.1.2030>

32. Sarkar, A., Xiaotong, Toronto, N., Drosos, I., & Poelitz, C. (2024). *When copilot becomes autopilot: Generative AI's critical risk to knowledge work and a critical solution*. arXiv. <https://doi.org/10.48550/arXiv.2412.15030>

33. Sherson, J. (2024). A hybrid intelligent change management approach to generative AI adoption. In *Frontiers in Artificial Intelligence and Applications*. <https://doi.org/10.3233/FAIA240228>

34. Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2024a). *Ironies of generative AI: Understanding and mitigating productivity loss in human-AI interactions*. arXiv. <https://doi.org/10.48550/arXiv.2402.11364>

35. Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2024b). Ironies of generative AI: Understanding and mitigating productivity loss in human-AI interactions. *International Journal of Human-Computer Interaction*, 1–18. <https://doi.org/10.1080/10447318.2024.2405782>

36. Singh, A., Taneja, K., Guan, Z., & Ghosh, A. (2025). *Protecting human cognition in the age of AI*. arXiv. <https://doi.org/10.48550/arXiv.2502.12447>

37. Smith, G., Luka, N., Lattimore, B. R., Newman, J., Nonnemeke, B., & Mittelstadt, B. (2025). Responsible generative AI use by product managers: Recoupling ethical principles and practices. *Academy of Management Proceedings*, 2025(1). <https://doi.org/10.5465/AMPROC.2025.24377a>

38. Tankelevitch, L., Glassman, E. L., He, J., Kittur, A., Lee, M., Palani, S., Sarkar, A., Ramos, G., Rogers, Y., & Subramonyam, H. (2025). *Understanding, protecting, and augmenting human cognition with generative AI: A synthesis of the CHI 2025 tools for thought workshop*. arXiv. <https://doi.org/10.48550/arXiv.2508.21036>

39. Tankelevitch, L., Kewenig, V., Simkute, A., Scott, A. E., Sarkar, A., Sellen, A., & Rintel, S. (2024). The metacognitive demands and opportunities of generative AI. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–18. <https://doi.org/10.1145/3613904.3642902>

40. Wang, W., Yang, M., & Sun, T. (2023). *Human-AI co-creation in product ideation: The dual view of quality and diversity*. SSRN. <https://doi.org/10.2139/ssrn.4668241>

41. Witkowski, A., & Wodecki, A. (2025). Where does AI play a major role in the new product development and product management process? *Management Review Quarterly*. <https://doi.org/10.1007/s11301-025-00533-5>

42. Yatani, K., Sramek, Z., & Yang, C.-L. (2024). *AI as extraherics: Fostering higher-order thinking skills in human-AI interaction*. arXiv. <https://doi.org/10.48550/arXiv.2409.09218>

43. Zhong, H. (2024). *Optimal integration: Human, machine, and generative AI*. SSRN. <https://doi.org/10.2139/ssrn.4899250>

44. Zysman, J., & Nitzberg, M. (2024). *Generative AI and the future of work: Augmentation or automation?* SSRN. <https://doi.org/10.2139/ssrn.4811728>

45. Agunbiade, O. L. (2025). Strategic Investment Analysis in Emerging Markets: A Framework for Value Creation, Financial Resilience, and Sustainable Private Equity Performance in Sub-Saharan Africa.