



OPEN ACCESS

SUBMITTED 21 October 2025

ACCEPTED 29 October 2025

PUBLISHED 12 November 2025

VOLUME Vol.07 Issue11 2025

CITATION

Hyo Ji Yun. (2025). The Influence of Artificial Intelligence on Paradigms of Entrepreneurship and Innovative Product Development. The American Journal of Interdisciplinary Innovations and Research, 7(11), 13–19.
<https://doi.org/10.37547/tajir/Volume07Issue11-02>

COPYRIGHT

© 2025 Original content from this work may be used under the terms of the creative common's attributes 4.0 License.

The Influence of Artificial Intelligence on Paradigms of Entrepreneurship and Innovative Product Development

Hyo Ji Yun

Co-Founder & CEO @ Givance
Raleigh-Durham-Chapel Hill Area, NC, US

Abstract: The study examines how artificial intelligence reshapes entrepreneurial activity and new-product trajectories across discovery, design, launch, and post-launch learning. The contribution lies in consolidating recent evidence on AI's impact on opportunity formation, decision flows, and governance in risk-sensitive and regulated domains such as mental health, education, and organizations and products across sectors. The review describes performance gains from generative systems in knowledge work, the structuring value of multi-agent orchestration, and the strategic implications for data provenance, evaluation discipline, and go-to-market execution. Special attention is paid to cross-phase integration, where insight artifacts travel intact from qualitative discovery to modeling and telemetry-guided iteration. The goal is to derive an operator's framework that aligns research signals with venture routines and product leadership choices. Methods include comparative synthesis of peer-reviewed sources, critical appraisal, and concept mapping. The conclusion outlines a reproducible stack for AI-native entrepreneurship that balances speed with auditability and trust, relevant for founders and product leaders.

Keywords: AI entrepreneurship, multi-agent systems, product data and user- and process-interaction signals, innovation management, product development, venture governance, generative AI, data provenance, causal evaluation, go-to-market

Introduction

poultry meat Artificial intelligence is altering the economics of search, the construction of evidence, and the cadence of decisions in new ventures. Market

sensing increasingly relies on large-scale embeddings and simulation; internal workflows rely on orchestrated agents that shorten iteration cycles; commercial execution relies on telemetry-updated assets rather than static plans. The study aims to systematize these changes and present an operator's framework suitable for AI-intensive product leadership in data-intensive markets with high-feedback loops.

The objectives are threefold:

- 1) to consolidate recent evidence on performance, limits, and transferability across venture stages;
- 2) to map multi-agent orchestration and data-ops to repeatable decision rituals in product management;
- 3) to translate governance concerns—provenance, consent, drift, evaluation—into actionable checkpoints for launch and learning.

Novelty stems from aligning heterogeneous findings into a stage-by-stage stack that binds discovery artifacts to modeling, governance, and go-to-market routines, with direct applicability to mental health, education, and growth and engagement analytics.

Materials and Methods

The evidence base draws on ten peer-reviewed sources from entrepreneurship, innovation, AI systems, management, education, and health. P. Davidsson [1] analyzes AI as an external enabler for new venture emergence. G. Giuggioli [2] reviews entrepreneurial adoption patterns and future research directions. X. Li [3] surveys LLM-based multi-agent workflows, tool usage, and orchestration challenges. M. M. Mariani [4] synthesizes AI in innovation research with a structured framework. M. Mariani [5] discusses generative AI implications for innovation management. S. Noy [6] reports randomized-trial evidence on productivity effects in knowledge work. F. Ouyang [7] reviews AI-driven learning analytics in collaborative learning environments. M. K. K. Rony [8] provides recent sectoral evidence on AI adoption and practitioner perceptions in clinical training. A. Witkowski [9] maps high-impact insertion points of AI across new product development and product management. Z. Zhang [10] examines AI-enabled business model innovation in digital platform enterprises.

A comparative literature synthesis was conducted, combining structured source analysis, concept mapping across venture stages, and critical appraisal of transferability to risk-sensitive and regulated domains.

The study applied qualitative content analysis, cross-study triangulation, and a stage-wise analytical framework to align findings with discovery, design, launch, and post-launch learning. The writing relied on comparative method, source analysis, evaluative synthesis, and conceptual integration.

Results

Rapid shifts appear in how entrepreneurs discover, validate and scale opportunities once artificial intelligence acts as an external enabler across the entire venture process, altering the salience of information signals, the cost of search, and the locus of expert judgment [1; 2]. Evidence from systematic reviews of innovation scholarship shows a durable transition from isolated, task-level analytics to end-to-end, learning systems that couple data generation with decision flows in product strategy [4; 5].

Founders working with data-rich problems register measurable gains once generative models compress iteration cycles in knowledge work. A large randomized trial with mid-career professionals reported ~40% lower task time and higher output ratings when generative AI assistance was allowed, a magnitude consistent with shorter loops between customer discovery and internal deliverables in early product work [6]. In product functions, these efficiency gains propagate unevenly: ideation, copy, lightweight research and draft analysis accelerate most, while evaluation quality depends on governance and prompt scaffolding, reshaping hiring, onboarding and review systems in entrepreneurial teams [5; 6].

Across new-product stages, recent synthesis identifies where AI contributes persistent leverage: opportunity sensing via market signals, usage telemetry, and operational events; concept screening through simulation and preference modeling; design-space search with generative architectures; demand forecasting and adaptive pricing for launch; and continuous post-launch optimization from telemetry and feedback. Fragmented, phase-specific deployments remain the modal pattern in the literature; the strongest gains occur when insight handoffs are engineered across phases rather than left as one-off tools [9]. Figure 1 summarizes these leverage points and highlights the cross-phase integrations most often missed in current practice.




Figure 1. AI leverage points across the new product development and product management process [9]

Multi-agent systems now provide a practical control layer for those cross-phase handoffs. A current survey of LLM-based multi-agent systems documents stable patterns—planner/critic/executor roles; tools for retrieval, orchestration and simulation; and protocols for task decomposition—that map cleanly to startup workflows: market signal harvesting, assumption tracking, experiment scheduling, and post-release triage [3]. In ventures that operate with thin managerial bandwidth, these agentic stacks formalize product rituals (PRD updates, decision logs, red-team passes), reduce tacit-knowledge loss during growth, and keep the data tagging → modeling → decision loop auditable under compliance constraints [3].

Opportunity formation, teaming and capability building shift in tandem. Entrepreneurship studies over the last few years converge on AI as an “enabler” that broadens who can found (lowering threshold skills for analysis and prototyping), expands the set of viable opportunities (through synthetic data and cheap scenario tests), and reweights human capital toward problem framing, governance and stakeholder sense-making [1; 2]. These effects manifest most clearly where problem definitions are data-intensive markets with high-feedback loops growth analytics, mental-health triage, or adaptive education—because signal curation, labeling strategies, and feedback design become the primary sources of differentiation rather than access to compute alone [2].

Sectoral evidence reinforces that pattern. In digital mental health, recent meta-analytic and review work catalogues clinically relevant uses—screening, risk

detection, relapse monitoring and patient-reported outcomes—while underscoring the fragility of generalization without rigorous cohorting, bias audits and human-in-the-loop escalation [8]. In education, a Q1 systematic review of AI-driven learning analytics in computer-supported collaborative learning reports consistent gains in engagement and feedback timeliness, paired with design cautions around transparency and instructor workload; those results map to venture products that promise adaptive guidance and affect-aware support, provided explainability and handoff affordances are engineered from the start [7].

Innovation-management research frames generative AI as a driver of new product trajectories and business architectures. Consolidated findings point to shifts in novelty search and concept diversity, new tensions in IP and data provenance, and ecosystem dynamics where model, data and distribution partners co-evolve with the product roadmap [4; 5]. For founders, that portfolio logic surfaces operationally as:

- 1) dedicated data-ops for labeling, consent and drift;
- 2) experimentation markets coupling synthetic users with causal evaluation;
- 3) governance gates that blend automated audits with expert review;
- 4) GTM assets that are updated by telemetry rather than campaign calendars [5; 10].

The overall picture that emerges across these literatures is a paradigmatic move from discrete tool adoption to continuously learning AI-driven product systems: opportunity sensing that fuses qualitative inquiry with large-scale embeddings; product engines that externalize tacit expert heuristics into reproducible agent workflows; and monetization loops that learn directly from user interaction streams. Where founders operationalize that stack—especially with multi-agent orchestration, careful data tagging, and staged evaluation—venture throughput improves without sacrificing the human judgment needed for sensitive domains such as mental health, education and social-impact services.

Discussion

A shift toward continuously learning, risk-sensitive and regulated product systems emerges once AI functions not as a point tool but as an external enabler that reorganizes search, inference, and judgment across the venture lifecycle; the pattern aligns with recent entrepreneurship syntheses and innovation reviews that document movement from manual opportunity scouting to data-driven discovery, scenario simulation, and telemetry-guided iteration that reshape how founders allocate attention and capital [1; 2; 4; 5; 9; 10]. For product leaders who rely on product data and user- and process-interaction signals—where data labeling, cohort definition, and feedback design govern the fidelity of downstream inferences—the literature situates differentiation less in raw compute and more in the architecture of signals, the reproducibility of decision flows, and the tight coupling between modeling and go-to-market operations [2; 4; 5; 9; 10].

The strongest operational leverage appears when cross-phase handoffs are engineered explicitly rather than left as fragmented tool adoptions. Reviews in innovation management and entrepreneurship describe repeatable gains once insight production (retrieval, tagging, synthesis) feeds directly into assumption tracking, experiment scheduling, and release governance, a structure echoed by recent mappings of AI's high-impact insertion points across new product development; the same sources warn that portfolio benefits degrade when telemetry and evaluation remain siloed from roadmap decisions [1; 4; 5; 9; 10]. Multi-agent orchestration now supplies a practical control layer for those handoffs: planner/critic/executor role patterns, tool-use protocols, and state-sharing channels stabilize product

rituals such as PRD updates, red-teaming, and decision logging without bloating headcount—an attractive property for ventures scaling in domains where human oversight and auditability matter [3].

Productivity experiments with generative systems clarify where acceleration materializes and where guardrails are needed. Controlled evidence shows large reductions in completion time and higher expert ratings for knowledge-work deliverables, yet performance lifts concentrate on tasks near the training distribution while declines appear on more complex boundary problems; for entrepreneurial teams this implies revised hiring and review systems, with scaffolding for prompt patterns, retrieval hygiene, and escalation criteria embedded into product operations rather than treated as one-off training. In practice, that means treating prompt libraries, retrieval templates, and critique prompts as versioned product assets that co-evolve with the codebase and data pipelines, with ownership seated jointly in product, data, and compliance functions [5; 6].

Founders operating in mental health, education, and B2B solutions under heterogeneous regulatory constraints face sharper stakes because target outcomes intertwine with human vulnerability and institutional trust. Reviews in education technology recommend transparency of analytics and workload-aware orchestration for instructors; psychiatry syntheses emphasize bias audits, cohort-specific validation, and human-in-the-loop escalation for risk-bearing decisions; innovation-management pieces add governance checkpoints for provenance, consent, and drift. Read together, these sources outline a deployment stance in which explainability and red-team procedures sit on the same cadence as feature sprints, with liability-aware handoffs baked into UX flows rather than appended late in the release cycle [4; 5; 7; 8].

The mapping of evidence to venture tasks in Table 1 frames how an AI-native product leader converts research signals into operating choices at each stage, building directly on the leverage points synthesized for new product development and product management in recent management scholarship and the systematizations summarized earlier in Figure 1 [4; 5; 9; 10].

Table 1 - Evidence-informed AI contributions across venture stages and their operating implications [1; 2; 4-6; 9; 10]

Venture stage	Predominant AI contribution (evidence basis)	Operating implication for founders/product leaders
Opportunity sensing	Large-scale text/image embeddings, market-signal harvesting, early persona discovery	Treat data collection and labeling as first-class backlog items; wire qualitative discovery to retrieval workflows so interviews refine embedding spaces.
Concept screening	Simulation, preference modeling, generative exploration of design spaces	Replace ad-hoc brainstorming with agent-mediated variant generation plus rubric-based critique prompts; archive critiques to train evaluators.
MVP build	Code, copy, and experiment-asset acceleration; automated PRD diffs	Gate perceived velocity with review rituals: senior approval on prompts/templates; pair automated diffing with decision logs for audit.
Launch/GTM	Demand forecasting, dynamic pricing/segmentation, channel message testing	Align marketing cadence to telemetry windows; replace calendar-only planning with experiment markets and rolling attribution models.
Post-launch learning	Telemetry fusion, anomaly and drift detection, causal evaluation	Standardize red-team drills and rollback criteria; schedule model refresh as a product event with owner, metrics, and risk narrative.

The table underscores a simple pattern: gains concentrate where insight artifacts travel intact between phases—interview notes into retrieval corpora, critique rubrics into agent prompts, telemetry into causal evaluators—mirroring the cross-phase logic distilled in the new-product process literature and the leverage points summarized earlier in Figure 1 [4; 5; 9].

Sector-specific constraints then shape governance choices. In mental health, meta-analytic and review findings converge on risk stratification, relapse monitoring, and patient-reported outcomes as promising targets balanced by strict escalation pathways and bias-aware validation. In collaborative learning,

gains in engagement and feedback timeliness appear when analytics close the loop to instructor intervention yet degrade when orchestration adds hidden workload or opaque scoring. Multi-agent surveys suggest composable role patterns for these governance demands—critics for harm scanning, planners for task decomposition, and auditors for provenance checks—while innovation reviews add IP and data-licensing frictions that influence partner strategy [3; 4; 5; 7; 8; 10]. Table 2 consolidates these domain-tuned controls for product leaders who operate across growth and engagement analytics, mental-health triage, and education.

Table 2 - Domain-tuned governance controls for AI-enabled products in sensitive markets [3-5; 7; 8; 10]

Domain focus	Recurrent risks in the literature	Governance and design controls
Digital mental health	Label bias, cohort shift, unsafe autonomy in triage	Human-in-the-loop escalation, cohort-specific validation, harm-scan critics in agent stacks, consent-aware data lineage

Collaborative learning tools	Opaque analytics, instructor overload, feedback quality variance	Transparent dashboards, workload-aware orchestration, rubrics that surface model rationales, staged automation with opt-outs
General innovation pipelines	IP/provenance uncertainty, partner lock-in, drift	Provenance tracking, model/data escrow clauses, scheduled drift audits, experiment marketplaces for roadmap decisions

The controls listed above align with an operator stance where measurement precedes scaling. For leaders focused on product data and user- and process-interaction signals, that stance translates into three recurring design moves grounded in the sources: keep labeling strategies and consent schemas version-controlled like code; instrument agent interactions so critique prompts, retrieval queries, and tool calls are inspectable; tie GTM decisions to causal evaluators rather than headline metrics alone [3; 4; 5; 7; 8; 10]. When these disciplines hold, the productivity lifts from generative assistance flow into durable venture throughput without hollowing out expert judgment, a balance particularly relevant to founders who build in mental health, education, and growth and engagement—markets in which value creation rides on trust and stakeholder legitimacy no less than on speed [5; 6; 7; 8].

Evidence on multi-agent systems further clarifies how small teams keep complexity tractable while scaling cross-phase learning. Planner–executor–critic triads map cleanly onto product discovery (planner frames hypotheses and success criteria), delivery (executor composes drafts, code, and experiment assets), and governance (critic conducts harm scans, bias checks, and provenance verification), with shared memory to avoid brittle handoffs; surveys of such systems collect patterns for tool choice, state sharing, and failure recovery that reduce tacit-knowledge loss as headcount grows [3]. In ventures that operate across heterogeneous stakeholders—patients and clinicians, teachers and students, donors and growth and engagement users and stakeholders—these agentic patterns stabilize meeting notes, decision logs, and release rationales into search-ready corpora that later power retrieval and onboarding [3; 5; 7; 8].

The entrepreneurship literature frames these operational shifts as a redistribution of founder skill portfolios. Signal curation, problem framing, and stakeholder sense-making rise in value while low-level drafting and exploratory analysis compress in time;

lower thresholds for prototyping broaden who can found, yet variance in outcomes widens unless teams institutionalize governance and measurement, a point made repeatedly across reviews and concept papers on AI-enabled venturing and innovation management [1; 2; 4; 5]. Productivity experiments remind operators that performance gains remain task-dependent; embedding prompt scaffolds, retrieval hygiene, and escalation rules into product operations avoids hidden performance cliffs when work drifts outside the familiar distribution [5; 6].

In strategy terms, business-model studies advise treating data partnerships, licensing, and distribution not as downstream concerns but as first-order design variables. Provenance, consent, and re-use determine the ceiling for model adaptation; partner contracts and escrow clauses shape bargaining power as providers evolve; telemetry-updated GTM assets supplant calendar-only campaigns in fast-learning products; these themes recur across innovation management and business-model research and sit naturally alongside the cross-phase leverage points summarized in Figure 1 [4; 5; 9; 10]. For AI-driven founders working on product data and user- and process-interaction signals, the same logic favors early investments in data-ops, experiment markets with synthetic users paired to causal evaluators, and legal architectures that keep future fine-tuning viable [4; 5; 10].

For an operator with a track record in multi-agent systems, adaptive learning, and growth and engagement analytics, alignment between literature and practice looks direct: wire qualitative customer discovery into retrieval corpora so operational signals travel into modeling; institutionalize agent critics for harm scans in mental-health pathways; publish evaluation narratives as part of releases; and couple GTM updates to telemetry rather than static calendars. The sources summarized here suggest that such a stack produces not only faster iteration loops but sturdier legitimacy in markets that tie adoption to

trust, precisely where product data and user- and process-interaction products aim to shift outcomes at scale.

Conclusion

The consolidation shows a stable pattern: measurable acceleration in ideation and draft production, reliable gains where agentic orchestration governs handoffs, and persistent risk where provenance, consent, bias control, and drift monitoring remain ad hoc. A durable stack for AI-native entrepreneurship binds qualitative discovery to retrieval corpora, formalizes planner–executor–critic cycles, and couples GTM decisions to causal evaluators with explicit rollback criteria. Data-ops, evaluation discipline, and contract architecture around models and data determine headroom for adaptation and bargaining power with partners. For product leaders working with product data and user- and process-interaction in mental health, education, and growth and engagement analytics, the framework yields faster iteration without eroding trust: consent-aware data lineage, cohort-specific validation, and red-team procedures run on the same cadence as feature sprints. The outlined approach answers the study's objectives by translating research signals into operating routines that maintain auditability while sustaining venture throughput, offering a practical guide for founders and AI-driven product teams.

References

- Davidsson, P., & Sufyan, M. (2023). What does AI think of AI as an external enabler (EE) of entrepreneurship? An assessment through and of the EE framework. *Journal of Business Venturing Insights*, 20, e00413. <https://doi.org/10.1016/j.jbvi.2023.e00413>
- Giuggioli, G., & Pellegrini, M. (2023). Artificial intelligence as an enabler for entrepreneurs: A systematic literature review and an agenda for future research. *International Journal of Entrepreneurial Behavior & Research*, 29(4), 816–837. <https://doi.org/10.1108/IJEBR-05-2021-0426>
- Li, X., Wang, S., Zeng, S., Liu, Y., Zhang, J., & Tang, J. (2024). A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges. *Vicinagearth*, 1, 9. <https://doi.org/10.1007/s44336-024-00009-2>
- Mariani, M. M., Machado, I., Magrelli, V., & Dwivedi, Y. K. (2023). Artificial intelligence in innovation research: A systematic review, conceptual framework, and future research directions. *Technovation*, 122(C). <https://doi.org/10.1016/j.technovation.2022.102623>
- Mariani, M., & Dwivedi, Y. K. (2024). Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research*, 175, Article 114542. <https://doi.org/10.1016/j.jbusres.2024.114542>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Ouyang, F., & Zhang, L. (2024). AI-driven learning analytics applications and tools in computer-supported collaborative learning: A systematic review. *Educational Research Review*, 44, 100616. <https://doi.org/10.1016/j.edurev.2024.100616>
- Rony, M. K. K., Ahmad, S., Das, D. C., Tanha, S. M., Deb, T. R., Akter, M. R., Khatun, M. A., Khalil, M. I., Peu, U. R., Parvin, M. R., Alrazeeni, D. M., & Akter, F. (2025). Nursing students' perspectives on integrating artificial intelligence into clinical practice and training: A qualitative descriptive study. *Health Science Reports*, 8(4), e70728. <https://doi.org/10.1002/hsr.2.70728>
- 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>
- Zhang, Z., Kang, Y., Lu, Y., & Li, P. (2025). The role of artificial intelligence in business model innovation of digital platform enterprises. *Systems*, 13(7), 507. <https://doi.org/10.3390/systems13070507>