

Agentic Product Evolution: An Empirical Study of How Product Strategy Must Evolve in the Age of Agentic AI

Suhas Rumalla¹

Senior Director of Product Management
Walmart, United States

John Felix²

Senior Product Manager
Walmart, India

ABSTRACT

Enterprises are aggressively pursuing Generative AI initiatives in today's highly competitive environment. Along this journey, many are attempting to move from simple AI features to autonomous systems (aka agentic AI, agentic systems, autonomous agents). Even after making substantial technology investments, most companies still struggle to move beyond pilots and proofs of concept, with many ultimately abandoning their agentic AI ambitions. The prevailing narrative continues to attribute these failures to insufficient technology maturity. Only a small percentage, roughly 10% have managed to implement agentic AI systems successfully at scale. While, technology readiness is certainly a factor, this paper argues that the core challenge lies in a broken product strategy that has not evolved to support autonomous agents. Most product strategies were never designed to incorporate agentic behaviour, shared decision-making or autonomous reasoning into systems. Based on the observations from 10 successful enterprise agentic systems implementations, this study proposes Agentic Product Evolution (APE), a strategy model tailored for agentic AI use cases. The framework is further validated through insights from more than 25 product leaders across global technology companies that have successfully scaled agentic AI capabilities. Early adopters of the APE witnessed a drop in strategic rework by an average of 36% and a boost in their stakeholders' confidence by up to 60% with average being 48%. This paper positions APE as a practitioner-centric playbook and research-based contribution to the nascent field of GenAI product strategy.

General Terms

Agentic System Design, Autonomous AI Agents, AI Product Strategy, Agent Operating Model, Generative AI adoption

Keywords

Enterprise AI, Agentic AI, Agentic Product Evolution(APE), Product Management, Human-AI Collaboration

1. INTRODUCTION

Enterprises are investing heavily in GenAI and the autonomous capabilities it promises, with leadership directions, ambitious technology roadmaps and accelerated AI pilots and POCs. Yet, many initiatives stall because teams approach agentic systems purely as a technological evolution while employing traditional product strategy constructs like static persona, outdated guardrails, usage-based success measurement. This creates friction among cross-functional teams, deadlocks in role clarity, slow progress and eventually lost momentum before delivering meaningful outcomes. The paper identifies strategic patterns from 10 enterprise case studies

that delivered significant results using agentic AI systems, validating the patterns against success stories across companies that have scaled agentic AI capabilities successfully [2, 3, 4]. The findings are presented as a practical, outcome-proven playbook to help organizations design and scale agentic systems with clarity, alignment, and measurable impact.

2. THEORETICAL FRAMEWORK

Enterprise technology adoption is rapidly shifting from generative AI to agentic systems, a trend highlighted in McKinsey's State of AI report, which notes growing interest in AI that can "act, decide, and operate autonomously" within business processes [5]. However, many organizations still view the GenAI wave solely as a technological disruption, prioritizing model integration and tooling over the strategic and organizational changes required [6]. Gartner's Emerging Tech Hype Cycle reinforces this, showing that AI experimentation is accelerating far faster than enterprise readiness [7]. PwC's Global CEO survey reveals that close to 50% of CEOs are integrating AI systematically into their technology platform while not even 25% do the same for their business strategy [8]. Figure 1 highlights the widening gap between how enterprises adopt AI at the technological front and how little they adapt their strategic direction to support it.

While technological evolution is certainly a necessity, previous research work emphasizes that realizing AI's full value requires both structured capability progression and systematic maturity building, as depicted in Figure 2 [1]. It also offers methods for evaluating and prioritizing GenAI initiatives, helping organizations identify which opportunities merit investment [1].

No existing literature addresses how product strategy itself must evolve to support agentic AI capabilities, especially for systems that make contextual decisions, share agency with humans, and require continuous governance. This gap underscores the need for a dedicated framework that guides organizations to establish the strategic foundations necessary for successful agentic AI implementation.

3. RESEARCH METHODOLOGY

3.1 Research Question

When agents are infused into software products, how should the existing product strategies evolve in enterprise settings?

3.2 Framework Derivation Approach

This research employs a two-phased qualitative, inductive, and interpretivist approach:

—**Phase 1 - Framework Formulation:** This phase involved thematic analysis of 10 successful enterprise agentic systems imple-

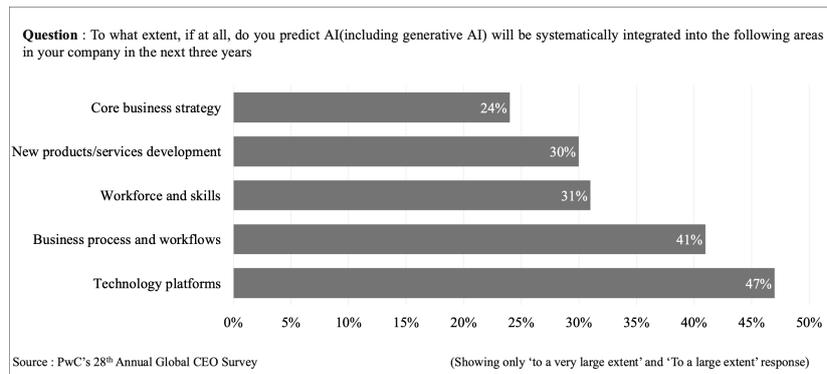


Fig. 1. CEO focus areas: comparison of AI integration into technology platforms versus business strategy

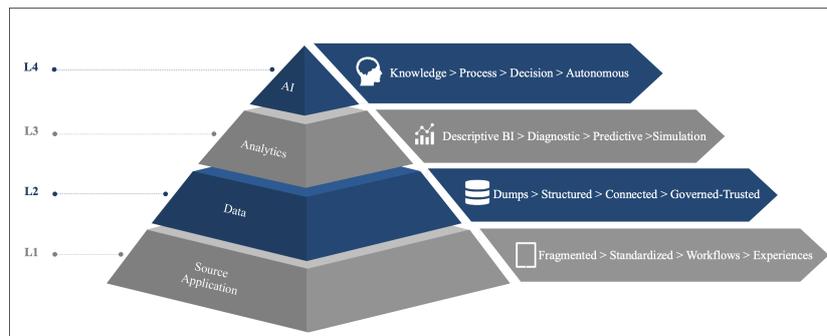


Fig. 2. Product maturity pyramid illustrating why each AI capability layer must be built sequentially to avoid stalled progress

mentations across diverse business functions. These cases served as structured success narratives, revealing how the enterprise re-defined its product strategies to support autonomous decision-making at scale. The insights captured from these implementations illustrate the strategic choices, operating models, and governance mechanisms required to operationalize agentic systems effectively within large enterprises.

Data Sources: The primary data consists solely of the documented case studies within the enterprise where agentic systems were delivering successful business results. No proprietary documents, or confidential artifacts were used. The cases serve as empirical evidence of evolved product strategy for agentic systems.

—**Phase 2 - Framework Validation:** This phase involved similar thematic analysis of public interviews, podcasts, and talks from various senior product leaders across global technology firms. This serves as a secondary validation, evaluating how closely the framework's (phase 1) patterns align with industry-wide thinking.

Data Sources: The secondary validation is based on transcribed interviews, podcasts, conference talks, and Q&As with over 25 product leaders from firms such as Microsoft, Google, Amazon, OpenAI, Meta, Anthropic, and others. Only the materials published in the last 24 months explicitly referencing AI integration, autonomy, or product strategy evolution were considered.

This two-phased approach uses case-based empirical evidence and thematic triangulation to create a product strategy framework that's

generic enough for the practitioners to manage their agentic AI initiatives effectively.

3.3 Data Analysis

3.3.1 Primary Case Analysis (Enterprise Case Studies). The 10 enterprise case studies were analyzed using a 3-stage inductive coding process [9, 10].

- Open Coding:** Identifying all meaningful observations from each case without imposing predefined categories (patterns)
- Axial Coding:** Grouping patterns into clear categories
- Selective Coding:** Defining the key themes from the categories, aligning to the primary case analysis

This analysis surfaced the strong themes across 10 case studies that explained the following within the enterprise: a) How agentic capabilities were embedded on different elements of product strategy, b) How decision autonomy, reasoning, trust, and guardrails were put into practice, and c) How organizational and architectural principles were transformed for agentic capabilities.

3.3.2 Secondary Thematic Validation (Cross-Industry). The same coding procedure was applied on the public interview transcripts of more than 25 product leaders from major global technology companies across industries [9, 10]. This analysis identified recurring themes in how industry leaders approached agentic AI. The resulting themes were then compared against those from the primary case analyses to validate whether the strategic patterns observed in the enterprise cases were also reflected across the broader industry.

3.3.3 Ethical Considerations

- The enterprise case studies used here are compiled success narratives and do not contain internal proprietary data or identifiable personnel information.
- All secondary sources are publicly available materials.
- No conflict of interest affects the integrity of the analysis.

4. RESULTS

Across the enterprise case studies, it became clear that niche business processes grew highly complex due to massive scale. These processes relied heavily on human judgment where single errors could ripple across thousands of transactions amplifying discrepancies, creating operational challenges, and ultimately imposing substantial financial burden on the enterprise. Workflow automation accelerated repetitive tasks but did little to improve decision quality in the workflows. It couldn't explain discrepancies, interpret ambiguous exceptions, detect anomalies dynamically, or adapt to external system changes. Leaders repeatedly observed that they had automated tasks, but had not strengthened decision-making in their systems [11, 12].

Product leaders recognized that agentic capabilities would fundamentally change how decisions were made inside their products [12, 13]. They proactively anticipated the need to rethink product strategy before introducing genuine decision-making actors [11]. They collaborated closely with their stakeholders to clarify decision ownership, define when autonomy was appropriate, and set expectations for oversight and trust. This early realization became a strategic inflection point that prompted them to reshape their strategy before deploying agentic systems; It ensured that a clear governance structure was in place to oversee any delegated decision-making. Figure 3 is the sample representation of a solution from the case studies that orchestrates the hybrid capabilities of the solution through a network of agents.

Below are the key elements in the product strategy that had been reimaged in the case studies for the successful implementation of agentic capabilities.

4.1 Vision - From Helping User with Tools to Human-Agent Decision Ownership

The vision evolved from streamlining fragmented workflows to enabling humans and agents to jointly ensure accurate and reliable operations. This required asking new questions: Which decisions must remain human-led? When can the agent act independently? What outcomes are shared? The reframed vision emphasized not just task-level efficiency but also continuous operational correctness. Setting the vision became about establishing a balanced partnership between human oversight and autonomous reasoning. This shift immediately reduced the strategic fog that had slowed teams for months.

4.2 Personas - From End-Users to a Multi-Actor Agent Ecosystem

Personas expanded from describing end-user tasks to defining a coordinated ecosystem of humans and agents, each with clear responsibilities, limitations, and escalation rules. In practice, this meant creating an AI counterpart to human roles - an "agent actor" empowered to handle routine tasks. Each agent was assigned a distinct decision function, such as interpreting discrepancies, detecting anomalies, guiding SOP steps, or adapting to changes in external systems. As a result, teams shifted from designing for the isolated

user experiences to designing for the interactions among all actors in the system.

4.3 Roadmap - From Feature Delivery to Autonomy Governance

Roadmaps, once centred on features and release dates, shifted to chart how autonomy would mature safely in terms of moving from insights to controlled automation, then to partial decision-making, and eventually to full delegation with oversight. This roadmap was no longer about shipping features; it was about calibrating the system behaviour. It emphasized teams to establish what agents were permitted to do, when autonomy could be expanded, and how risks would be governed throughout the progression of autonomous decision-making in the systems.

4.4 Architecture - From Service Integration to Cognitive Infrastructure

Architecture expanded beyond integrating services to facilitate reasoning, memory, observability, and adaptability as core capabilities. A cognitive layer in the architecture became essential for interpreting signals, correlating events, and acting consistently across diverse operational conditions. Explainability and auditability of the autonomous actions, were required to ensure safety and trust. The system ultimately evolved into an infrastructure built for making real-time decisions rather than simple processing transactions.

4.5 Adoption - From User Training for Use to Building Trust for Delegation

The User adoption strategy also shifted: the goal was no longer about training the users on a UI, but enabling them to build trust on the system by helping them understand why an agent acted a certain way. This adoption strategy was primarily focussed on giving the users, the confidence to delegate decisions without feeling a loss of control. Overall, the adoption strategy was defined by the users' confidence in delegation rather than proficiency with the tool.

4.6 Metrics - From Activity Tracking to Intelligence Performance

The success metrics shifted from tracking throughput and errors to evaluating how effectively the agents reasoned and made decisions. Most of the key indicators included decision accuracy, autonomy rate, override behaviour, false-blocks patterns, and self-healing actions. These metrics served as early warning signals, showing when autonomy was drifting or when reasoning required correction. The success of the system was measured by consistent, safe, and explainable agent decisions, rather than by system usage.

5. DISCUSSION

5.1 Agentic Product Evolution (APE) Framework

The shifts seen in the earlier section show a clear pattern: once agentic capabilities are introduced, product strategy needs to evolve. The following key elements under product strategy, i.e. vision, personas, roadmap, architecture, adoption, and metrics should undergo changes to redefine decision roles, collaboration methods, and accountability within the product. How the key element of product strategy transforms when agentic capabilities are introduced is depicted in Figure 4.

Building on these shifts, this paper defines Agentic Product Evolution (APE), a framework that guides leaders to create and man-

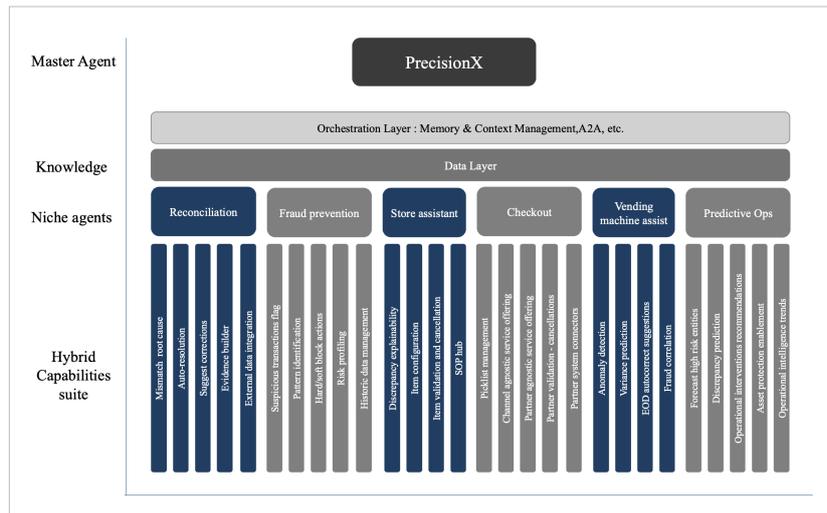


Fig. 3. Example of multi-agent orchestration demonstrating how hybrid human-agent capabilities jointly operate within an enterprise workflow

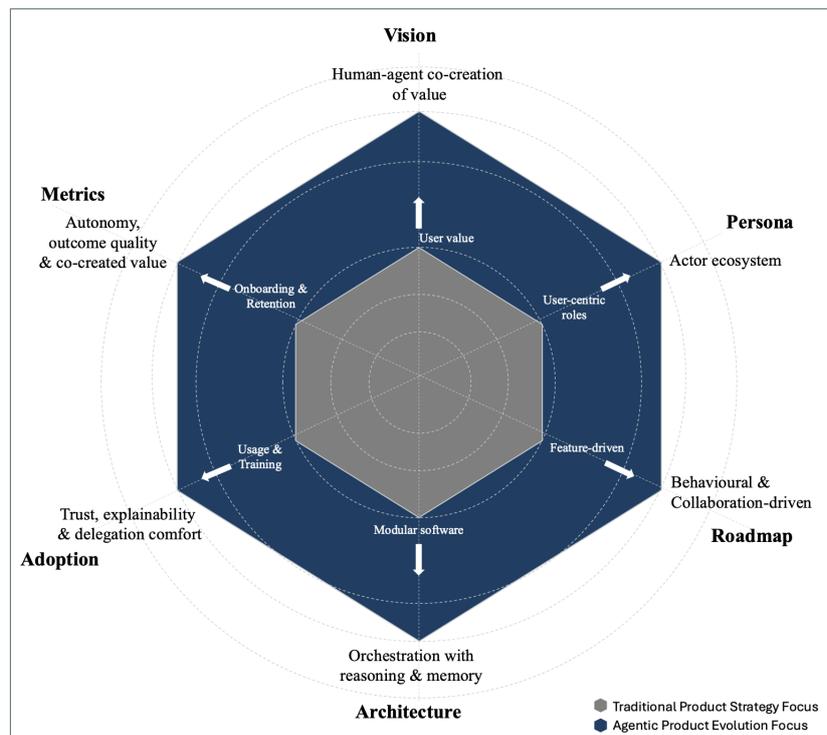


Fig. 4. How each product-strategy element evolves when agentic capabilities are introduced under the Agentic Product Evolution (APE) framework

age products where autonomous agents work alongside people. The call-to-action (CTA) for the leaders on each strategic elements' evolution is highlighted in Figure 5.

5.2 Cross-Industry Thematic Validation

To assess whether the APE framework aligns with broader industry practice, a thematic analysis was conducted on more than 25 public interviews with senior product leaders from Meta, Ope-

nAI, Amazon, Google, Microsoft, Apple, Atlassian, and others [14, 15, 16, 11, 3, 18, 19, 17, 21, 22, 23, 24, 25, 26, 27]. The objective was to evaluate how organizations implementing agentic AI systems successfully, reflect the strategic shifts proposed in the APE framework. This analysis examined whether their real-world approaches converge around the core principles outlined in the APE framework.

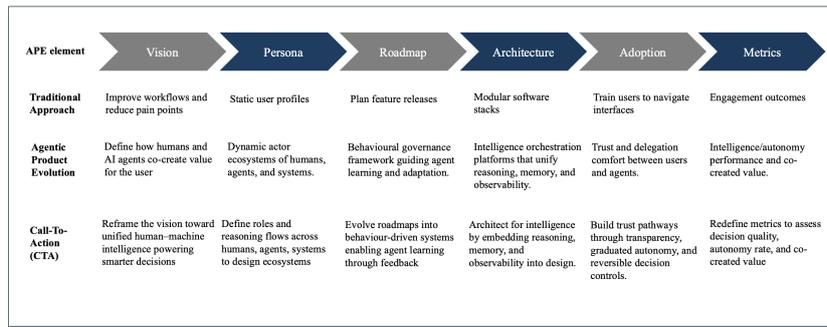


Fig. 5. Strategic call-to-actions for leaders: actions required across each product-strategy element to successfully implement agentic systems

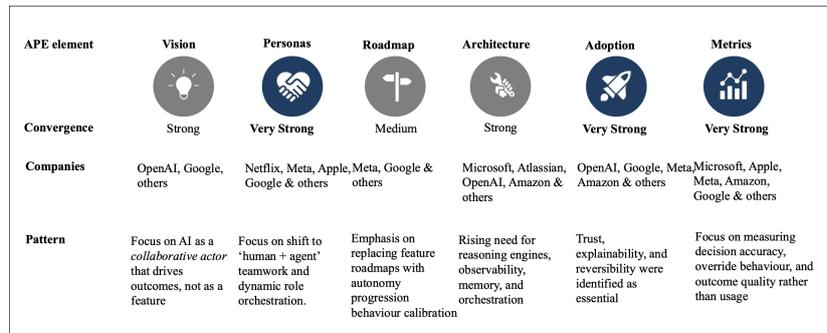


Fig. 6. Alignment scores showing how strongly global product leaders' perspectives converge with the strategic shifts defined in the APE framework

To score each element in the APE, the interviews were thematically coded first, to capture what different product leaders emphasized [9, 10]. Each interview was then evaluated on how often they mentioned a theme (frequency), how many interviews reflected the same idea (breadth), and how strongly or explicitly they expressed it (intensity). Based on these signals, each theme was tagged using clear thresholds: Very Strong (85-100) when the theme was widespread, cross-company, and expressed with high clarity; Strong (70-84) when it appeared frequently but with some maturity differences; Medium (60-69) when alignment was partial; and Low (45-59) when only a few leaders mentioned it.

The validation summary is presented in Figure 6. In a nutshell, the product leaders across different companies and industries, regularly noted the same shifts on the strategy elements that were framed in the APE. Their insights support the idea of APE being a new industry-agnostic strategic model for creating, managing, and expanding agentic AI systems.

6. IMPLICATIONS

With the APE framework, the enterprise was able to unlock immediate and measurable strategic lift in their business processes.

—First, the product strategic rework dropped sharply by 40% because teams finally had clear roles, decision boundaries, and shared understanding of how the system should work. This eliminated the back-and-forth clarification cycles that once slowed execution among the cross-functional teams; It also improved the speed of decision-making across functions.

—Second, product-level risk management was improved with the framework providing a structured approach to control autonomy.

This ensured safe decision-making and prevented unexpected agent behaviour.

—Third, internal Net Promoter Scores (NPS) increased by 20% due to boosted stakeholder confidence, which was influenced by the framework supporting a) transparent progression of autonomy, b) user-trust oriented metrics definition and c) robust governance and guardrails.

In total, Agentic Product Evolution (APE) provides a scalable and future-proof model - one that could enable enterprises to expand agentic capabilities into other areas without reinventing the strategy every time.

The table below summarizes the dos and don'ts for product leaders aiming to implement agentic AI capabilities in their system.

7. CONCLUSION

This paper presents the Agentic Product Evolution (APE), a practitioner-focussed model illustrating how product strategy must evolve when autonomous agents become active participants in enterprise systems. It reframes the core strategic elements (vision, personas, roadmap, architecture, adoption, and metrics) around a main goal of orchestrating shared human-agent intelligence.

For executives, the implications are very clear. Product leaders must move from feature thinking to orchestration of decision autonomy. Engineering leaders must design architectures with built-in reasoning, memory, and safe autonomy. Business leaders must establish adoption models grounded in trust and responsible delegation.

For practitioners and organizations, APE serves as a blueprint for scaling agentic systems safely and predictably. For researchers, it extends product management thinking into the era of autonomous decision-making.

Table 1. Key behaviours product leaders should adopt—and avoid—when designing and scaling agentic AI systems

S.No.	Dos	Don'ts
1	Set clear boundaries for decisions between humans and agents, including escalation rules	Don't view autonomy as a simple feature - it requires a complete redesign rather than small updates
2	Design for trust, ensuring actions can be explained, reversed, and involve confidence indicators	Don't use agents without clearly defined roles or safety measures - uncertainty creates operational risks
3	Develop systems with built-in reasoning, memory, observability, and adaptability	Don't try to increase autonomy faster than trust growth - user confidence drives adoption, not training volume
4	Implement ongoing governance through performance metrics and real-world signals like overrides and anomalies	Don't depend on traditional usage metrics - they can obscure issues with intelligence quality and safety

Products are no longer static tools, they are intelligent collaborators. Organizations that embrace this shift meticulously will stand a chance to pioneer the next era of digital innovation.

8. LIMITATIONS AND FUTURE RESEARCH DIRECTIONS

The study is based on qualitative data, providing a conceptual overview of how product strategy changes in the agentic era. Its dependence on sources from major technology organizations limits contextual depth and empirical validation. Future research should include: a) detailed case studies with quantitative evaluations to assess real-world application and impact of APE, b) wider industry and regional comparisons to test its significance beyond large tech companies, and c) more thorough investigation of the ethical aspects of autonomy, accountability, and bias.

9. REFERENCES

- [1] S. Rumalla, "Strategic decision-making framework for evaluating and selecting GenAI use cases," *International Journal of Science and Research Archive*, vol. 16, no. 01, pp. 170-177, 2025.
- [2] M. Purdy, "What is Agentic AI, and How will it change work," *Harvard Business Review*, 12 December 2024. [Online]. Available: <https://hbr.org/2024/12/what-is-agentic-ai-and-how-will-it-change-work>.
- [3] K. Ng, "HubSpot SVP of Product on AI Agents to Augment GTM Teams," *Product School*, 05 March 2025. [Online]. Available: <https://productschool.com/resources/product-podcast/karen-ng-hubspot-ai>.
- [4] H. B. Review, "Driving Business Value With Agentic AI," *Harvard Business Review*, 2025.
- [5] M. report, "The state of AI in 2025 - Agents, innovation, and transformation," *QuantumBlack AI by McKinsey*, 2025.
- [6] Deloitte, "Deloitte's State of Generative AI quarterly survey," *Deloitte Insights*, 2025.
- [7] Gartner, "2025 Gartner Hype Cycle for Emerging Technologies," *Gartner Industry report*, 2025.
- [8] PwC, "PwC's 28th Annual Global CEO Survey," *ceosurvey.pwc*, 2025.
- [9] A. L. Strauss and J. M. Corbin, *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*, Sage Publications, Inc., 1998.
- [10] V. & C. V. Braun, "Using thematic analysis in psychology – methodological grounding for your thematic analysis of interview transcripts," *Qualitative Research in Psychology*, vol. 3, pp. 77-101, 2006.
- [11] C. N. Desk, "Innovating Retail: Walmart Global Tech's Vision and the Future of Shopping," *CXO today.com*, 17 May 2024. [Online]. Available: <https://cxotoday.com/interviews/innovating-retail-walmart-global-techs-vision-and-the-future-of-shopping>.
- [12] L. Wilkinson, "Inside Walmart's all-in playbook for generative AI, agentic tools," 24 April 2025. [Online]. Available: <https://www.ciodive.com/news/walmart-agentic-generative-ai-strategy-developer-experience/745820/>.
- [13] B. Njuguna, "Walmart bets on AI agents to rewire developer workflows," *Silicon Angle*, 30 September 2025. [Online]. Available: <https://siliconangle.com/2025/09/30/walmart-leverages-developer-focused-ai-agents-boost-software-innovation-aifactoriesdatacenters/>.
- [14] S. Goldman, "Salesforce AI CEO Clara Shih says AI is a 'moving target' — but her aim is steady," *Venturebeat.com*, 21 01 2024. [Online]. Available: <https://venturebeat.com/ai/salesforce-ai-ceo-clara-shih-says-ai-is-a-moving-target-but-her-aim-is-steady>.
- [15] Salesforce+, "A Conversation with Sundar Pichai and Marc Benioff," *Salesforce*, 23 September 2025. [Online]. Available: www.salesforce.com/plus/experience/dreamforce_2025/.
- [16] K. Yeung, "Microsoft's AI chief: Why AI should be emotionally intelligent and how it'll create the next 1,000 Einsteins," *Venturebeat*, 21 May 2024. [Online]. Available: <https://venturebeat.com/ai/microsoft-ai-chief-mustafasuleyman-build-2024-interview>.
- [17] P. Yang, "What Survives vs. Gets Disrupted by AI," *Behind the Craft*, 31 August 2025. [Online]. Available: <https://creators.spotify.com/pod/profile/peter-yang42/>.
- [18] M. Marshall, "How Intuit killed the chatbot crutch – and built an agentic AI playbook you can copy," *Venturebeat*, 29 August 2025. [Online]. Available: <https://venturebeat.com/ai/how-intuit-killed-the-chatbot-crutch-and-built-an-agentic-ai-playbook-you-can-copy>.
- [19] C. G. d. Villaumbrosia, "Intercom CPO on shifting from AI-dead to AI-first," *The Product Podcast*, 15 January 2025. [Online]. Available: <https://theproductpodcast.buzzsprout.com/90361/episodes/16400006/>.
- [20] P. Yang, "How to Get Your Team to Start Prototyping with AI," *Behind the Craft*, 03 August 2025. [Online]. Available: <https://createconomy.so/p/how-to-get-your-team-to-start-prototyping-with-ai-jeremy-vanta>.
- [21] D. Shipper, "Transcript: 'Microsoft's AI Vision: An Open Internet Made for Agents'," *Every.to*, 20 May 2025. [Online]. Available: <https://every.to/podcast/transcript-microsoft-s-ai-vision-an-open-internet-made-for-agents>.
- [22] A. Gupta, "Rubrik CPO on Scaling AI, Cybersecurity & Business Model Transformation," *Product School*, 17 September 2025. [Online].

- [23] H. Langley, "Google just gave us a tantalizing glimpse into the future of AI agents," *Business Insider*, 14 May 2024. [Online].
- [24] J. DeLanghe, "Slack's VP of Product & GM Reveals How to Monetize AI Product Features," *Product School*, 09 August 2024. [Online].
- [25] B. L. & L. Downes, "Can AI Agents be Trusted?" *Harvard Business Review*, 26 May 2025. [Online].
- [26] C. G. d. Villaumbrosia, "Twilio CPO on Integrating AI into Product Strategy to Grow Revenue," *The Product Podcast*, 27 August 2025. [Online].
- [27] L. Columbus, "Top ten ways Intuit is revolutionizing personalization with generative AI," *Venturebeat*, 15 July 2024. [Online].
- [28] C. Christensen, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business Review Press, 2024.
- [29] R. G. Cooper, *Winning at New products: Creating Value Through Innovation*, Hachette Audio, 2023.
- [30] M. & Company, "The state of AI: How organizations are rewiring to capture value," *QuantumBlack AI by McKinsey*, 12 March 2025. [Online].
- [31] S. Ventures, "Salesforce Ventures Deploys Over \$850M of Its \$1B AI Fund," *Salesforce Ventures*, 16 October 2025. [Online].
- [32] J. L. & G. Crossan, "Autonomous generative AI agents: Under development," *Deloitte*, 19 November 2024. [Online].