Operationalizing Generative AI in Software Product Management: A Review of Managerial Use-Cases, Governance, and Ethical Guardrails

Gaurij Mahajan
Email: mahajangaurij@gmail.com

***Abstract*— This paper synthesizes recent evidence on how generative AI reshapes software product management across discovery-to-delivery workflows, emphasizing managerial decisions, outcomes, and governance. Drawing on studies spanning market analysis, positioning, customer insight, requirements engineering, Agile execution, UI/UX, and engineering productivity, the paper maps concrete applications to established product management domains and highlights observable effects on efficiency, quality, and customer experience. Using change and strategy lenses, the analysis outlines adoption prerequisites—strategy alignment, role design, process integration, data readiness, and risk controls—alongside an ethics agenda covering bias, privacy, accountability, and IP exposure. The contribution distills a practical blueprint for product leaders: where to deploy generative AI for business impact, how to embed it within portfolio and lifecycle decisions, and which guardrails enable responsible scaling. The primary contribution of this paper is a novel conceptual framework that integrates GenAI capabilities into the ISPMA lifecycle, providing a structured model for adoption, governance, and impact assessment for both researchers and practitioners.**

***Index Terms*—Generative AI, software product management, ISPMA, governance, ethics, Agile.**

I. Introduction

Generative AI, residing at the cusp of the second and third wave of AI evolution as classified by DARPA (2023), is a testament to the transformative potential of AI. While the first wave comprised rule-based AI systems, and the second wave advanced to machine learning and deep learning applications, the third wave anticipates AI systems capable of contextual adaptation, akin to human understanding. Straddling the second and third phases, generative AI leverages statistical learning to generate novel content such as text, images, or music, often utilizing deep learning frameworks such as Generative Adversarial Networks (GANs). As we progress into this intriguing realm, we unlock vast possibilities for technological transformation.

Generative AI refers to artificial intelligence that can create new content, such as images, text, or music, which did not exist before. It often involves learning from a large amount of input data and then generating something new that reflects the patterns or features observed in that data. In recent years, Artificial Intelligence Generated Content (AIGC) has not only intrigued the computer science community but also captivated societal interest with various products developed by major tech firms, including OpenAI's conversational AI system, ChatGPT, and their creative visual AI, DALL-E2. These models, underpinned by deep generative AI using neural networks, have made significant strides over the past decade [1]. The most famous form of generative AI is the generative adversarial network (GAN), which consists of two neural networks competing to generate suitable content material. Other generative AI strategies include variational autoencoders (VAE) and transformers, including OpenAI's GPT-3, demonstrating unique competencies in producing human-like textual content [2].

Software product management is a crucial discipline that encompasses the activities and responsibilities involved in creating, delivering, and maintaining software products. Following the guidelines set by the International Software Product Management Association (ISPMA), software product managers are equipped with the knowledge and skills to effectively navigate the complex landscape of product strategy, planning, development, and management. By employing a holistic approach, software product management aims to align business objectives with customer needs, optimize product features, prioritize requirements, coordinate cross-functional teams, and oversee the entire product lifecycle. With its emphasis on user-centricity, value creation, and Agile methodologies, ISPMA provides a comprehensive framework that empowers software product managers to drive innovation, foster collaboration, and achieve sustainable success in today's dynamic software industry [3].

Software product management is a complicated and challenging discipline that includes various activities: market analysis, product analysis, product strategy, product planning, development, marketing, sales, support, and services. The achievement of software product management relies upon the software product manager's ability to understand the users' needs, market trends, and technological improvements. Generative artificial intelligence (GenAI) has emerged as a promising technology that may help software product managers enhance their decision-making and the quality of their products. This systematic literature evaluation aims to offer an overview of the current state of research on the applications of generative AI in software product management, focusing on the software product manager's activities. The potential applications of generative AI in software product management are vast, and this article offers an in-depth evaluation of how these applications can transform the software product

management field. The following sections will discuss the literature review and the benefits and potential limitations of incorporating generative AI in software product management.

A. Problem Statement

Generative AI is a revolutionary technology with the potential to modernize software product management by automating tasks, improving efficiency, and enhancing customer experience [4]–[6]. Growing demand to modernize workflow across industries is expected to drive the demand for generative AI applications among industries. The global market for generative AI is expected to reach $109 billion by 2030 [7]. As per the McKinsey & Company (2022) report, a 67% average share of respondents reported a revenue increase via AI adoption and 79% reported a cost decrease via AI adoption. The general problem is the need for software product managers to access new and innovative tools that can help them improve their decision-making and profitability and enhance the quality of their products. The specific problem is the lack of awareness of generative AI's potential in software product management. This study aims to bridge this knowledge gap, creating a more comprehensive understanding of how generative AI can be utilized in software product management along with the ethical considerations. It will empower product managers to leverage this technology, potentially leading to the creation of superior products.

II. Literature Review

This systematic literature review is based on a comprehensive analysis of existing studies related to the applications of generative AI in business and product management. The ISPMA (2023) framework is used to study applications in product management. The International Software Product Management Association (ISPMA) Product Management Framework is a well-known model designed specifically for managing software products. Paajoki (2020) identifies the ISPMA Product Management Framework as the best practice framework to adopt for any organization. The aim of this review is to identify and synthesize the current state of research, focusing on key application areas such as market research, product planning, product documentation, product requirements engineering, product development, UI/UX design, customer insights, and Agile software development. The McKinsey 7-S Framework and Lewin's Change Model frame the review to evaluate how an organization can adopt this technology change, laying the foundation for product managers and business leaders to use the technology effectively and efficiently [8], [9]. Owen et al. (2013)'s Responsible Innovation framework and General Data Protection Regulation (GDPR, 2023) principles frame the discussion on ethics and privacy.

A. Applications of Generative AI in Software Product Management

1) Idea Generation and Market Research: Karim et al. (2022) researched the application of generative AI for idea generation, brainstorming, and producing research clues. Their work builds on prior research that developed a model for generating medical dialogues related to COVID-19. The authors utilized two GPT-3-based models: GPT-NEO-125M and GPT-NEO-1.3B. The larger GPT-NEO-1.3B model consistently generated more coherent and interconnected ideas. Although their study focused on the medical domain, specifically COVID-19, they suggest that the principles derived from their findings could be extended to other scientific or specialized disciplines. The findings show that the larger model generates more coherent text with linked ideas.

Large Language Models (LLMs), such as GPT, are a type of AI created to understand and generate human-like language. Brand et al. (2023)'s research demonstrated that these models can be instrumental in understanding customer preferences. They found that, when treated like a randomly selected customer, GPT exhibits responses that are realistic and consistent with values obtained from existing research. Although they employ traditional market research paradigms to underscore the utility of GPT, LLMs may inspire new market research paradigms unrestricted by the constraints of human subject research. The authors also cautioned that LLMs are known to occasionally "hallucinate" and return incorrect information.

2) Customer Insights and Support: Siggelkow and Terwiesch (2023) discussed the transformative potential of large language models such as OpenAI's ChatGPT and Google's Bard in enhancing customer experiences. They highlighted that these AI models can aid in recognizing customer needs by interpreting and integrating data, translating these needs into specific requests, and responding to customers with tailored solutions. They emphasized the importance of focusing on the customer, not just the technology, by addressing specific customer pain points and deploying AI to complement, not replace, a firm's unique capabilities.

Additionally, Siggelkow and Terwiesch (2023) underscored the value of the "repeat" dimension in customer interactions, where firms learn from each interaction to improve future responses. Generative AI systems excel in this area, creating a positive feedback loop that enhances AI's understanding of the customer for more personalized and effective service. The authors warned against viewing AI as a substitute for human labor, suggesting instead that it should be seen as a tool to enhance a firm's capabilities in unique ways. Ultimately, while generative AI has the potential to revolutionize customer experiences, the authors emphasized the importance of strategic deployment to address customer needs and bolster a firm's distinctive value proposition.

Brynjolfsson et al. (2023) investigated the implementation of a generative AI conversational assistant across 5,179 customer support agents, finding that the AI tool notably increased worker productivity, enhanced customer sentiment, and decreased employee turnover rates. Particularly beneficial for newer and less-skilled workers, the AI tool substantially improved problem resolution and customer satisfaction but did not significantly assist the highest-skilled or more experienced workers. An analysis of the text from agent conversations sug-

gested that AI recommendations guided less-skilled workers to communicate more similarly to high-skilled workers.

3) Product Requirements Engineering and Agile Development: Malik et al. (2022) focused on the critical process of Requirement Engineering (RE), which is integral to software development and involves defining, documenting, and maintaining software requirements. They noted the significance of Software Requirement Specifications (SRS), key deliverables in the software development life cycle, and the problems arising from ambiguities and conflicts within these documents. To tackle this issue, they proposed a two-phase process for automatic conflict detection in SRS documents that works directly with natural language requirements. The first phase involves transforming software requirements into numeric vectors using transformer-based sentence embeddings, with cosine similarity and ROC curves employed to identify potential conflicts. In the second phase, potential conflicts undergo analysis with general and software-specific Named Entity Recognition (NER), and an overlapping entity ratio is used to determine the final set of conflicts. The authors found an improvement in the F1-score for the OpenCoss and IBM-UAV datasets by 4% and 5%, respectively, underlining its effectiveness.

In Agile software development, generative AI can support various activities, such as sprint planning, backlog management, and estimation. For instance, Kim et al. (2021) proposed an AI-driven approach for estimating user story points in Agile projects, leveraging natural language processing techniques to analyze and estimate the effort required for user stories. Fu and Tantithamthavorn (2022) introduced GPT2SP, an innovative Agile story point estimation approach, utilizing a GPT-2 pre-trained language model and Transformer-based architecture. The proposed method surpasses traditional techniques, such as Planning Poker, Analogy, and expert judgment, and Deep-SE, a deep learning-based approach. Through testing on over 23,000 issues spanning 16 open-source projects, GPT2SP exhibited superior accuracy, outperforming within-project estimates by 34%–57%, cross-project estimates by 39%–49%, and enhancing Deep-SE's performance by 6%–47%. A proof-of-concept tool was also developed to clarify the factors influencing the estimates. A practitioner survey highlighted the challenge of story point estimation. However, it also indicated that an AI-based approach with explanations, such as GPT2SP, was deemed more valuable and trustworthy. Dam et al. (2019) introduced a framework that adapts and integrates various AI technologies to bolster Agile project management. Although areas in Agile project management are still challenging due to insufficient adequate support, their proposed analytics engine aims to provide decision support on multiple fronts. This includes descriptive analytics, a basic level of analytics that most existing Agile project management tools offer, primarily through data visualization via reports, dashboards, and scorecards.

4) Automated Code Generation and UI/UX Design: There have been several types of studies on code generation using generative AI techniques such as Transformer, but Peng et al. (2023) found that GitHub Copilot (which uses generative AI technology for code generation), an AI pair programmer providing context-based code suggestions, significantly boosts productivity. The study revealed that the group using Copilot completed tasks 55.8% faster, marking the first experiment of its kind to provide empirical evidence of the potential of AI tools to enhance human productivity. If extrapolated, a 55.8% productivity increase could lead to substantial economic cost savings and notably influence GDP growth.

Khan and Uddin (2022) examined the application of the Generative Pre-trained Transformer-3 (GPT-3) Codex in automating documentation generation in Software Engineering (SE). Codex demonstrates state-of-the-art performance, outperforming prior models by achieving an average BLEU score of 20.63, even under basic settings. Unlike previous approaches that required task- or language-specific retraining or fine-tuning, Codex operated efficiently with one-shot learning, a process explored by randomly selecting one sample from the corresponding training set.

Park et al. (2023) discussed developing and evaluating the ALSI-Transformer, a transformer-based code comment generation model designed to improve source code comprehension in software development. In response to the increasing need for efficient code commenting as the scale of open-source software grows, the ALSI-Transformer uses a novel method of aggregating multimodal information through Gate Network.

The integration of AI technology is revolutionizing the Human-Computer Interaction (HCI) and User Experience (UX) landscapes, redefining user research, design, and evaluation methodologies [10]. AI-based solutions enhance the UX quality, with increasing adoption prompted by heightened user awareness of technical innovation. Houde et al. (2022) explored the objectives, obstacles, and operational practices of teams working on application modernization, as showcased by three projects. The authors discovered that user experience (UX) modernization is a complex, labor-intensive segment of the broader modernization process, involving numerous manual tasks performed by multiple team members, such as project managers, UX designers, and software engineers. Despite the crucial role of UX in these projects, this is often inadequately addressed compared to other tasks such as core code transformation into microservices. However, the authors recognized its potential as a promising avenue for implementing generative AI technologies. By identifying specific pain points in the UX modernization process, they proposed using generative AI models as a possible solution. They also envisioned a future scenario where these models could redefine the UX modernization workflow.

B. Ethical Considerations in Generative AI

Brand et al. (2023) highlighted the problem of "hallucinations" and returning incorrect information with Large Language Models. It requires consistent human reviews to ensure the accuracy of the AI-generated output. The rise in deep generative AI models has led to increasingly intricate models. However, their performance relies heavily on quality training data. Notably, these models often exhibit the "black

box" problem, limiting their interpretability and causing potential trust issues. Lastly, handling AI-generated content must address critical social concerns to ensure responsible and beneficial usage for society, emphasizing the importance of trustworthiness and responsibility in the field [1].

Generative AI presents legal risks, including potential infringement of intellectual property rights, including unresolved legal questions such as ownership and application of copyright, patent, and trademark laws to AI-generated content. Before leveraging the benefits of generative AI, businesses must understand these risks and devise strategies for self-protection. Essential measures include updating vendor and customer contracts with explicit disclosure about the use of generative AI and clauses safeguarding intellectual property rights. Furthermore, confidentiality provisions should be enhanced to prohibit confidential information in AI tool text prompts [11].

Dwivedi et al. (2023) discussed several key challenges that the rise of generative AI, such as ChatGPT, presents to the industry. They highlight serious ethical concerns, as these AI models lack the ability to understand or consider ethical and legal issues. This leads to potential misuse, including the production of deepfakes and disinformation. The inherent "black box" nature of AI systems poses transparency and explainability issues. Biases can be replicated from the training data, unintentionally leading to misinformation. Legal issues abound, with few guidelines for AI development, unclear copyright boundaries, and dubious ownership of AI-generated content. The authors also foresee potential job losses and increased technology dependency following AI incorporation, stunting personal development and leaving organizations vulnerable if technology fails. Additionally, generative AI systems are limited to combining existing information, offering limited originality. The cultural and personal acceptance of AI could lead to a new form of digital divide. Lastly, the effective use of generative AI requires the ability to design efficient prompts, potentially necessitating widespread new skill training.

III. Research Methodology

This study employs a Systematic Literature Review (SLR) following the guidelines proposed by Kitchenham and Charters [**?**] to identify, evaluate, and synthesize the existing research on Generative AI in Software Product Management.

A. Research Questions

The SLR was guided by the following research questions:

- **RQ1:** What are the reported use-cases and applications of Generative AI across the software product management lifecycle?
- **RQ2:** What are the proposed governance mechanisms and ethical guardrails for using Generative AI in SPM?
- **RQ3:** What are the identified gaps and future research directions in this domain?

B. Search Strategy

A comprehensive search was conducted across five major digital libraries: IEEE Xplore, ACM Digital Library, Google Scholar, EBSCOhost, and ProQuest Central. The search string was designed using a population-intervention-context (PIC) structure:

```
("generative AI" OR "large language model" OR "LLM"
    OR "GPT" OR "ChatGPT")
  AND ("software product management" OR "product
 manager" OR "requirements engineering" OR "Agile
      development" OR "UX design")
```

C. Study Selection and Inclusion Criteria

The study selection process followed the PRISMA flow diagram [**?**], as shown in Figure 1. The inclusion and exclusion criteria are detailed in Table I.

TABLE I
Study Inclusion and Exclusion Criteria

Inclusion Criteria	Exclusion Criteria
Peer-reviewed journal articles, conference papers, and workshop papers.	Books, editorials, prefaces, and non-peer-reviewed magazine articles.
Studies published between 2019 and 2023.	Studies published before 2019.
Primary focus on GenAI application in SPM or closely related fields (RE, Agile, UX).	Studies where AI/ML is mentioned but not generative AI specifically.
Written in English.	Non-English publications.
Full text available.	Abstracts only with no full text.

D. Data Extraction and Synthesis

A total of 78 primary studies were selected for final analysis. Data was extracted into a spreadsheet, capturing: publication venue, year, research method, GenAI technology, SPM activity, and key findings. The ISPMA framework [**?**] was used as a lens for thematic synthesis, categorizing the findings into the domains outlined in Section II.

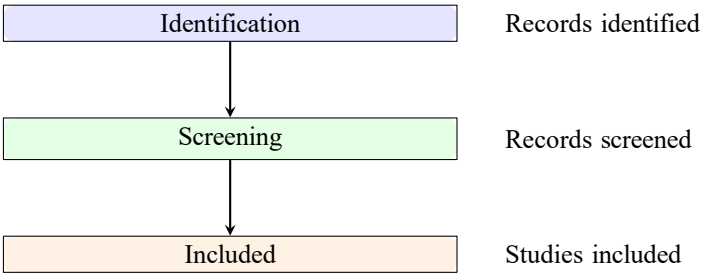

Fig. 1. PRISMA Flow Diagram illustrating the study selection process.

IV. A Framework for Operationalizing GenAI in SPM

Based on the synthesis of literature in Section II, we propose a conceptual framework for operationalizing Generative AI in Software Product Management. The framework, illustrated in Figure 2, is designed to help product leaders strategically integrate GenAI capabilities across the product lifecycle while implementing necessary governance structures.

TABLE II
SUMMARY OF GENERATIVE AI APPLICATIONS IN SOFTWARE PRODUCT MANAGEMENT

SPM Domain (ISPMA)	GenAI Applications & Use-Cases	Reported Benefits & Outcomes
Market Research & Strategy	Acting as a synthetic customer for surveys; generating market trends analysis reports.	Faster idea validation; uncovering non-obvious customer segments.
Product Requirements Engineering	Automated conflict detection in SRS; generating user stories from epics.	Improved requirement quality; reduced manual review effort.
Agile Development & Planning	AI-based story point estimation (e.g., GPT2SP); sprint planning assistance.	Higher estimation accuracy; increased team velocity.
UI/UX Design	Generating design prototypes from text prompts; automating usability feedback analysis.	Rapid prototyping; more user-centric design iterations.
Customer Insights & Support	Analyzing large volumes of feedback; powering conversational support agents.	Enhanced customer sentiment; increased agent productivity.

A. Framework Components

The framework consists of three interconnected layers:

1) **SPM Lifecycle Layer (ISPMA-Based):** This layer represents the core product management activities, structured according to the ISPMA framework [**?**]. It spans from strategic (Product Strategy, Planning) to tactical and operational activities (Development, Launch, Monitoring).
2) **GenAI Capability Layer:** This layer details the specific GenAI functionalities that can be applied to enhance the SPM activities. These include Idea Generation, Content Synthesis, Automated Code Generation, and Predictive Estimation, among others.
3) **Governance & Guardrails Layer:** This foundational layer encompasses the essential structures required for responsible and effective adoption. It includes Ethical Guidelines (e.g., bias mitigation, fairness), Data Governance (privacy, quality), and Organizational Readiness (strategy, skills, change management).

B. Framework Dynamics

The framework posits that successful operationalization is not merely applying a GenAI tool to a single task. Instead, it requires:

- **Vertical Alignment:** A conscious mapping of specific GenAI capabilities (Layer 2) to address pain points or create leverage in specific SPM phases (Layer 1). For example, using LLMs for market research (Capability) directly supports Product Strategy (SPM Phase).
- **Horizontal Governance:** Every application of a GenAI capability must be underpinned by the corresponding guardrails (Layer 3). An AI used for customer insight analysis must adhere to data privacy rules, while one used for code generation must have IP safeguards.

This framework provides a structured approach for practitioners to move from isolated experiments to a coherent, responsible, and scalable GenAI strategy within SPM.

C. Critical Analysis and Tensions

While the potential benefits are significant, our synthesis reveals several critical tensions that product managers must navigate:

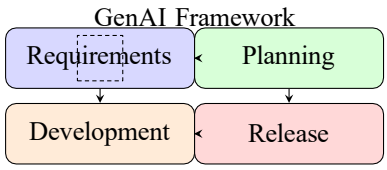

Fig. 2. GenAI in Software Product Management

Automation vs. Human Judgment: GenAI can automate tasks like user story writing or code generation. However, over-reliance risks *deskilling* and a loss of critical contextual judgment. The role of the product manager may shift from creator to curator and validator of AI-generated content.
Speed vs. Accuracy & Hallucination: The dramatic increases in productivity (e.g., 55.8% faster task completion with GitHub Copilot [4]) are compelling. Yet, this speed is contingent on rigorous human review to mitigate the risk of AI "hallucinations" and the propagation of incorrect or biased information.
Personalization vs. Data Privacy: Using GenAI to create hyper-personalized customer experiences requires deep data access, which directly conflicts with stringent data privacy regulations like GDPR. Organizations must find a balance, potentially through synthetic data generation or advanced anonymization techniques.

These tensions underscore that operationalizing GenAI is not a purely technical challenge but a socio-technical one, requiring careful attention to process redesign and role definition.

V. RESULTS AND ANALYSIS

In software product management, generative AI has demonstrated potential across various applications, such as market analysis, positioning and product definition, customer insights and support, product requirements engineering, and development execution. Studies show that large language models, such as GPT-3, can generate more coherent and interconnected ideas, create visually appealing and unique user interfaces, automatically estimate user story points, elaborate on natural language code, and analyze customer feedback [12], [13]. Additionally, generative AI can automate the process of requirement elicitation and analysis and software documentation, helping managers better understand and manage product

requirements while reducing manual effort. Furthermore, generative AI has the potential to revolutionize market research by using large language models [14].

Overall, the applications of generative AI in software product management are vast and diverse, ranging from product discovery to delivery. By leveraging the capabilities of generative AI models, software product managers can automate tedious tasks, make data-driven decisions, and improve overall product quality. The integration of generative AI in software product management processes can potentially lead to better product outcomes, streamlined workflows, and more efficient use of resources across various stages of product development [6], [12], [13], [15]–[17].

A. Applications of Generative AI in Software Product Management

The applications of generative AI in software product management can be categorized as follows:

Market Research: Generative AI models such as GPT can be used in market research, acting as hypothetical customers to understand preferences. Models such as these could inspire new research paradigms unrestricted by human subject research. Generative AI can generate ideas and brainstorm in specialized disciplines, with larger models generating more coherent and interconnected ideas [12], [14].

Positioning and Product Definition: GenAI, primarily used for content writing, can aid in product positioning and definition, including naming and feature suggestions, with demonstrated efficacy in writing e-commerce product descriptions and title generation across various industries [18], [19].

Customer Insights and Support: Generative AI models such as ChatGPT and Google's Bard can enhance customer experiences by interpreting and integrating data, translating needs into requests, and responding with tailored solutions. They can also analyze large volumes of customer feedback data to extract valuable insights. Additionally, generative AI can create a positive feedback loop, enhancing an understanding of the customer for more personalized service [6].

Product Requirements Engineering: Generative AI can automate requirement elicitation and analysis, generating natural language requirement statements from structured data. Also, it can help detect conflicts in Software Requirement Specifications (SRS) documents [5].

Agile Development: Generative AI can support sprint planning, backlog management, and estimation in Agile software development. AI models such as GPT-3 can automatically generate user stories based on provided input data, saving time, reducing human error, and improving the quality of user stories [13], [16].

Automated Code Generation: Generative AI can aid in automated code generation, enhancing productivity and efficiency. Tools such as GitHub Copilot have demonstrated a significant boost in productivity. Furthermore, generative AI can be used to automate documentation generation and improve the comprehension of source code in software development [4], [20], [21].

UI/UX Design: Generative AI can create visually appealing and unique interfaces, streamlining the design process and enabling rapid prototyping. It can generate UI/UX designs based on user preferences and requirements, leading to more user-centric designs and improved product experiences [15].

Business Decision-Making: Generative AI can enhance decision-making in business by addressing limitations of the bounded rationality model, with younger users and males typically more open to this technology's adoption [22].

B. Ethical Implications and Mitigation Strategies

The ethical implications of generative AI (GenAI) are diverse and consequential, underlining the need for rigorous guidelines and regulatory measures to ensure its responsible usage [1], [11], [14], [23].

Fairness: Bias is an inherent risk in GenAI. The AI system can inadvertently incorporate and propagate biases present in the data used for training, leading to skewed or prejudiced outputs [14]. Therefore, it is essential to ensure fairness and diverse training data to avoid potential harm or discrimination.

Data Privacy: Data privacy is another critical ethical consideration. As GenAI systems require extensive data for training, the likelihood of infringing on data privacy increases. Adhering to privacy regulations such as the GDPR (2023) is a step towards ensuring data privacy. GDPR's principles provide a comprehensive roadmap for data collection, storage, and usage [24], guiding businesses in the ethical use of personal data for GenAI.

Accountability: Accountability in GenAI involves taking responsibility for the system's outputs. If an AI system generates harmful or inaccurate content, the onus is on the organization to rectify the situation and prevent recurrence [1].

Transparency: Potential misuse of GenAI-produced content is another pressing concern. There is the risk of fabricated or manipulated content being passed off as legitimate, which could have severe consequences on fields such as news and journalism, and on financial reports and legal documents [11]. Rigorous verification mechanisms are necessary to prevent such misuse. The increase in deep generative AI models has led to increasingly intricate models. However, their performance relies heavily on quality training data. Notably, these models often exhibit the "black box" problem, limiting their interpretability and causing potential trust issues [1].

Robustness: Brand et al. (2023) highlighted the problem of "hallucinations" and returning incorrect information with large language models. This requires consistent human reviews to ensure the accuracy of the AI-generated output.

Legal Risks: Generative AI presents legal risks, including potential infringement of intellectual property rights, including unresolved legal questions such as ownership and application of copyright, patent, and trademark laws to AI-generated content. Before leveraging the benefits of generative AI, businesses must understand these risks and devise strategies for self-protection. Essential measures include updating vendor and customer contracts with explicit disclosure about the use of generative AI and clauses safeguarding intellectual

property rights. Furthermore, confidentiality provisions should be enhanced to prohibit confidential information in AI tool text prompts [11].

The Responsible Innovation (RI) framework, proposed by Owen et al. (2013), offers a comprehensive approach to addressing these ethical issues. The framework encourages anticipation, reflexivity, inclusion, and responsiveness to ensure the balanced development and application of GenAI. The GDPR and RI principles, when appropriately applied, can help mitigate these ethical challenges and guide the responsible use of GenAI. However, regulatory challenges persist. Hacker et al. (2023) argued that the current EU AI Act fails to keep pace with the rapid advancements and impact of Large Generative AI Models (LGAIMs). The authors advocated for a more technology-neutral approach to regulation, highlighting the need to maintain online discourse civility, create a level playing field for new AI models, and keep up with rapid advancements in AI technologies.

VI. Insights for Practitioners

Based on our synthesis and the proposed framework, we distill the following actionable insights for software product managers and leaders seeking to operationalize Generative AI:

- **Start with Strategy, Not Technology:** Align GenAI initiatives with specific product goals and pain points. Use the framework in Section IV to identify where GenAI can create the most leverage, rather than adopting tools indiscriminately.
- **Prioritize Governance from Day One:** Establish clear protocols for data privacy, IP protection, and output validation *before* scaling use-cases. The "Horizontal Governance" layer is non-negotiable for responsible adoption.
- **Augment, Do Not Automate:** The most effective use of GenAI augments human decision-making. Use AI for data synthesis and idea generation, but retain human judgment for strategic decisions, customer empathy, and validating AI-generated content against business context.
- **Invest in Prompt Engineering Skills:** The quality of GenAI output is highly dependent on input. Developing skills in crafting effective prompts is a critical new competency for product teams.
- **Measure Impact Rigorously:** Track metrics beyond productivity (e.g., task speed). Focus on outcome-oriented KPIs such as requirement quality, customer satisfaction scores, and product success rates to truly assess GenAI's value.

These insights emphasize that successful operationalization is a socio-technical challenge, requiring changes in process, skills, and strategy alongside technology adoption.

VII. Conclusion and Future Work

This study synthesized the current landscape of Generative AI applications in Software Product Management through a systematic review of 78 primary studies. By mapping use cases to the established ISPMA framework and analyzing associated benefits and challenges, we provided a structured overview for both researchers and practitioners. Furthermore, we proposed a conceptual framework (Section IV) to guide the integrated and responsible adoption of GenAI, emphasizing the alignment of capabilities with SPM activities and the foundational role of governance.

A. Limitations

A primary limitation of this review is its reliance on emerging, yet often preliminary, academic and industry reports. The field is evolving rapidly, and longitudinal studies on the sustained impact of GenAI on business KPIs (e.g., product success rates, long-term team productivity) are still scarce. Furthermore, our focus was on identifying applications, and a meta-analysis of the quantitative effect sizes was beyond our scope.

B. Future Research Directions

Based on our analysis, we propose the following concrete research directions:

1) **Empirical Validation of the Framework:** Conduct case studies or action research to validate and refine the proposed GenAI-SPM integration framework in real-world product organizations.
2) **Causal Impact on Quality:** Move beyond correlation to establish causal inference. For example, controlled experiments are needed to determine if AI-generated user stories actually lead to fewer defects or higher customer satisfaction compared to human-written ones.
3) **Governance Mechanism Efficacy:** Investigate which specific governance mechanisms (e.g., specific prompt review protocols, IP clauses) are most effective at mitigating risks like bias, IP infringement, and hallucinations in an SPM context.
4) **Human-AI Collaboration Models:** Explore optimal interaction patterns between product managers and AI assistants. This includes designing AI systems that enhance, rather than replace, strategic human judgment and creativity.

In conclusion, Generative AI holds transformative potential for SPM, but realizing this potential requires a deliberate, research-informed approach that balances innovation with responsibility. We hope this review and the proposed framework serve as a foundation for future work in this dynamic area.

References

[1] Y. Cao, S. Li, Y. Liu, Z. Yan, Y. Dai, P. S. Yu, and L. Sun, "A comprehensive survey of AI-generated content (AIGC): A history of generative AI from GAN to ChatGPT," *arXiv preprint* arXiv:2303.04226, 2023.

[2] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Scaling laws for autoregressive generative modeling," *arXiv preprint* arXiv:2102.08602, 2021.

[3] ISPMA, "A comprehensive guide for effective software product management," 2023. [Online]. Available: https://ispma.org/ook/

[4] S. Peng, E. Kalliamvakou, P. Cihon, and M. Demirer, "The impact of AI on developer productivity: Evidence from GitHub Copilot," *arXiv preprint* arXiv:2302.06590, 2023.

[5] G. Malik, M. Cevik, D. Parikh, and A. Basar, "Identifying the requirement conflicts in SRS documents using transformer-based sentence embeddings," *arXiv preprint* arXiv:2206.13690, 2022.

[6] N. Siggelkow and C. Terwiesch, "Create winning customer experiences with generative AI," *Harvard Business Review*, 2023. [Online]. Available: https://www.hbr.org/2023/04/create-winning-customer-experiences-with-generative-ai

[7] Grand View Research, "Generative AI market size to reach $109.37 billion by 2030," 2023. [Online]. Available: https://www.grandviewresearch.com/press-release/global-generative-ai-market

[8] T. J. Peters and R. H. Waterman, *In Search of Excellence*. Harper & Row, 1984.

[9] K. Lewin, "Frontiers in group dynamics: Concept, method and reality in social science; social equilibria and social change," *Human Relations*, vol. 1, no. 1, pp. 5–41, 1947.

[10] W. Xu, "AI in HCI design and user experience," *arXiv preprint* arXiv:2301.00987, 2023.

[11] G. Appel, J. Neelbauer, and D. A. Schweidel, "Generative AI has an intellectual property problem," *Harvard Business Review*, 2023. [Online]. Available: https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem

[12] M. R. Karim, S. S. Antar, and M. A. Khan, "Idea generation using transformer decoder models," in *Proceedings of the 2022 5th International Conference on Algorithms, Computing and Artificial Intelligence*, 2022, pp. 1–5.

[13] H. Kim, B. Choi, and H. Cho, "AI-driven user story points estimation based on natural language processing," *Journal of Systems and Software*, vol. 171, p. 110818, 2021.

[14] J. Brand, A. Israeli, and D. Ngwe, "Using GPT for market research," *SSRN*, no. 4395751, 2023.

[15] S. Houde, S. I. Ross, M. Muller, M. Agarwal, F. Martinez, J. Richards, and J. D. Weisz, "Opportunities for generative AI in UX modernization," in *Joint International Conference on Intelligent User Interfaces Workshops*, 2022.

[16] M. Fu and C. Tantithamthavorn, "GPT2SP: A transformer-based agile story point estimation approach," *IEEE Transactions on Software Engineering*, 2022.

[17] H. K. Dam, T. Tran, J. Grundy, A. Ghose, and Y. Kamei, "Towards effective AI-powered agile project management," in *2019 IEEE/ACM 41st International Conference on Software Engineering: New Ideas and Emerging Results (ICSE-NIER)*, 2019, pp. 41–44.

[18] M. T. Nguyen, P. T. Nguyen, V. V. Nguyen, and Q. M. Nguyen, "Generating product description with generative pre-trained transformer 2," in *2021 6th International Conference on Innovative Technology in Intelligent System and Industrial Applications (CITISIA)*, 2021, pp. 1–7.

[19] M. Zhang, Z. Gang, W. Yu, N. Huang, and W. Liu, "MaaPTG: Multimodal aspect-aware product title generation," *Journal of Intelligent Information Systems*, vol. 59, no. 1, pp. 213–235, 2022.

[20] J. Y. Khan and G. Uddin, "Automatic code documentation generation using GPT-3," in *37th IEEE/ACM International Conference on Automated Software Engineering*, 2022, pp. 1–6.

[21] Y. Park, A. Park, and C. Kim, "ALSI-Transformer: Transformer-based code comment generation with aligned lexical and syntactic information," *IEEE Access*, 2023.

[22] P. Korzynski, G. Mazurek, A. Altmann, J. Ejdys, R. Kazlauskaite, J. Paliszkiewicz, and E. Ziemba, "Generative artificial intelligence as a new context for management theories: Analysis of ChatGPT," *Central European Management Journal*, 2023.

[23] Y. K. Dwivedi, N. Kshetri, L. Hughes, E. L. Slade, A. Jeyaraj, A. K. Kar, and R. Wright, "'So what if ChatGPT wrote it?' Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy," *International Journal of Information Management*, vol. 71, p. 102642, 2023.

[24] GDPR, "Principles relating to processing of personal data," 2023. [Online]. Available: `https://gdpr-info.eu/art-5-gdpr/`