

Supporting Problem Exploration in Early-Stage Innovation Projects through Human-AI Collaboration

Anonymous Author(s)

Abstract

Context: Innovation is an important driver of economic growth, yet it often fails in its early stages because teams miss real market needs. AI is increasingly introduced in innovation, mostly in product development rather than in understanding real customer needs. *Objective:* To address this gap, this research explores how human-AI collaboration can assist in early-stage innovation, such as ideation and problem exploration. *Method:* This study adopts the three-cycle view of design science research. We ensured relevance through field observations of innovation teams at a design sprint event and derived their challenges. We then developed an initial prototype concept to address those challenges and created an evaluation plan to test the prototype in a realistic context. *Results:* We found that most teams fail in the earliest stages of innovation by missing customer-problem fit, that is, by failing to identify a meaningful problem for a real customer. AI only partially helped, allowing teams to externalise thoughts but not to deepen their understanding of the problem. We therefore designed a prototype that leverages empathy through storytelling to create a more realistic understanding of customers. *Next Steps:* The immediate next step is to implement the prototype and evaluate it in a realistic context to inform the next iterations of this study.

CCS Concepts

• **Human-centered computing** → **Collaborative interaction**; • **Computing methodologies** → *Artificial intelligence*; • **Software and its engineering** → Software creation and management.

Keywords

early-stage innovation, design sprint, human-AI collaboration, customer-problem fit, design science research

ACM Reference Format:

Anonymous Author(s). 2026. Supporting Problem Exploration in Early-Stage Innovation Projects through Human-AI Collaboration. In *Proceedings of The 30th International Conference on Evaluation and Assessment in Software Engineering (EASE 2026)*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

1 Introduction

Innovation is a key driver of economic growth, with R&D investment, high-tech exports, and scientific publications all shown to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EASE 2026, Glasgow, United Kingdom

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
<https://doi.org/10.1145/nnnnnnnn.nnnnnnn>

causally predict GDP growth in developing economies [13]. Yet, many innovation efforts fail because teams do not address real market needs¹, and early-stage startup teams in particular are prone to building solutions to problems that do not exist [4].

Generative AI (GenAI) could, in principle, address this challenge. Capabilities most relevant to early-stage innovation, including synthesizing customer feedback, analyzing market signals, and mapping competitive landscapes, have been identified as potential use cases of large language models (LLMs) [16]. In practice, however, AI adoption in innovation has concentrated almost entirely in development activities [16]. In software innovation, for instance, GenAI is increasingly being used to accelerate prototype development through human-AI collaboration, including emerging approaches such as vibe coding [1, 3, 17].

While GenAI and human-AI collaboration are accelerating solution development, the early phase of innovation, including ideation and problem exploration, remains comparatively slow and error-prone. As a result, teams may simply get faster at building the wrong thing. This research therefore explores whether human-AI collaborative techniques similar to vibe coding can be adapted to the underexplored area of early-stage innovation to improve problem exploration and understanding, and ultimately produce more aligned, better-grounded solutions. This motivates the following leading research question:

RQ: How can human-AI collaboration support early-stage innovation projects?

To answer this question, this research adopts a design science research approach [6], specifically the three-cycle view [5], comprising the relevance, design, and rigor cycles. As this study is still in an early stage, this paper reports primarily on the relevance cycle and outlines a plan for the design cycle.

Within the *relevance cycle*, we conducted a field observation at a design sprint event where novice innovation teams (student teams from diverse backgrounds) tackled innovation challenges across multiple domains. Despite the guidance provided by the design sprint method, most teams struggled to reach customer-problem fit, staying at a high level of abstraction and failing to identify a concrete, meaningful problem to solve for a specific customer. We also found that GenAI use reinforced this pattern, acting as a “comfort blanket” that reduced perceived uncertainty without pushing teams toward more concrete problems. These findings informed the *design cycle*, in which we defined an artefact concept designed to address the identified challenges and an evaluation plan, while implementation and final empirical validation are planned as subsequent work.

The rest of the paper is structured as follows. Section 2 presents the theoretical background on the fuzzy front end, customer-problem-solution fit, and design sprints. Section 3 describes the design science research approach and the field observation setup. Section 4

¹<https://www.cbinsights.com/research/report/startup-failure-reasons-top/>

reports the key challenges observed during the design sprint event. Section 5 derives design implications for problem framing and AI collaboration. Section 6 introduces the artefact concept and its three phases. Section 7 outlines the planned longitudinal evaluation plan. Section 8 discusses links to design thinking and hypothesis-driven development. Finally, Section 9 summarizes contributions and next steps.

2 Background

2.1 Fuzzy Front End of Innovation

The fuzzy front end (FFE) is the earliest phase of innovation, occurring before formal development begins, when teams explore opportunities and shape initial concepts [7, 11]. It is described as “fuzzy” because goals, user needs, and solution directions are ambiguous and volatile, and teams must work with incomplete and evolving information [11].

Despite this ambiguity, the FFE is particularly consequential: decisions and assumptions formed at this stage propagate into later phases, making them difficult and costly to reverse [7, 8]. Therefore, existing work has proposed methods, tools, and techniques to be used in the FFE. For example, Koen et al. [10] structure the FFE using the NCD model (New Concept Development, in contrast to New Product Development), and propose that teams use different tools depending on uncertainty and context, instead of a single standard process. It is important to build a common language for front-end activities, make assumptions explicit early, and use structured methods to reduce ambiguity before formal development. Among various practical tool examples suggested, lead-user and ethnographic/customer-insight methods are most relevant for customer and problem exploration.

2.2 Customer-Problem-Solution Fit

A typical goal of the early-stage of innovation is to reach customer–problem–solution (CPS) fit, which is the state in which a team has identified a specific type of customer who experiences a real and significant problem, and has developed a solution concept that addresses that problem in a way that customers find compelling [14]. As a result, to reach CPS fit, all three components must align: a real problem for the wrong customer, or the right customer but a solution they do not value, does not result in a complete fit [12, 14].

Reaching CPS fit is important because it prevents teams from refining solutions in search of a problem. Without it, efforts can drift toward improving features for customers who do not experience the problem or addressing issues that are not meaningful enough to drive adoption. This wastes development effort and can lock teams into early assumptions that are hard to reverse.

2.3 Design Sprint

The design sprint is a well-established process initially defined by Google Ventures [9] aiming to find CPS-fit. It builds on techniques such as design thinking and is typically performed by a small, cross-functional team in a time-boxed manner over five days, where each day focuses on a specific part of early-stage innovation:

- (1) **Understand:** The problem is mapped and explored, and a specific challenge to focus on is chosen.

- (2) **Sketch:** For the chosen problem, teams ideate a range of solutions, exploring a broad set of possibilities.
- (3) **Decide:** One specific solution is chosen from the ideas generated, often through structured voting techniques such as dot voting.
- (4) **Prototype:** The chosen solution is prototyped to produce a testable and realistic version of the idea.
- (5) **Test:** The prototype is tested with real customers to obtain actionable insights.

For each day, design sprints provide dedicated toolkits and techniques. On the *Understand* day, tools such as empathy maps, user interviews, personas, lightning talks, and How Might We notes are commonly used to build a shared understanding of the problem space. On the *Sketch* day, techniques like Crazy 8s (asking teams to develop eight ideas in eight minutes) and structured solution sketches are employed to generate and capture a broad range of ideas. On the *Decide* day, methods such as dot voting, the heat map critique, and storyboarding help teams converge on a single solution. On the *Prototype* day, teams typically produce low- to mid-fidelity prototypes using tools such as paper prototyping or digital mockups, guided by the storyboard produced the previous day. Finally, on the *Test* day, structured user interviews and think-aloud protocols are used to evaluate the prototype and gather actionable insights. While the original format spans five days, adapted shorter variants exist in practice [9].

Despite methods like design sprints, many innovation efforts still fail by building solutions without a clear problem to solve. This study examines why these failures persist and explores how human–AI collaboration, which has shown promise in improving development efforts, can be leveraged at this early stage to better surface genuine user needs.

3 Research Method

This study adopts the design science research (DSR) methodology proposed by Hevner et al. [6]. It specifically adopts the three-cycle view, comprising the relevance, design, and rigor cycle [5]:

- **Relevance cycle:** The relevance cycle connects the research to real-world needs by identifying and validating the practical problem to address. In line with this goal, we conducted field observation at a design sprint event with nine multidisciplinary student teams (54 students in total), where most were novice innovators. Each team tackled one of three open-ended challenges (see Table 1) provided by industry partners. This setting was chosen deliberately, because the design sprint methodology is a well-established technique for guiding early-stage innovation, making it a strong context for studying where early-stage innovation still fails even with formal processes in place.
- **Design cycle:** The design cycle focuses on designing, implementing, iterating, and evaluating an artefact that addresses the identified needs. To this end, we developed an initial artefact concept based on the problems identified in the field study. The artefact initially targets novice innovation teams, the same user group observed in the relevance cycle,

as they are most likely to face challenges in early-stage innovation.

We also developed an evaluation plan for the prototype, which will be conducted with novice innovation teams. The plan specifies the evaluation method and criteria to generate insights for future research iterations.

- **Rigor cycle:** The rigor cycle links the work to prior knowledge and develops meta-artefacts for scientific contribution and validity. For this, we based our design concept on established techniques such as design thinking as well as lean startup. As meta-artefacts we aim to develop human-AI collaborative approaches to help novice teams navigate the early, fuzzy stage of innovation.

4 Field Observations

This section reports the challenges we observed in the first run of the relevance cycle during the design sprint event.

4.1 Staying at the Surface of the Problem

We consistently observed that teams treated the challenge provided to them as the final problem to solve, rather than translating it into a concrete customer need. For instance, a team working on C3 (increasing the engagement of younger generations in cooperative bank membership), approached it as an organisational process to fix. They never asked why a specific young person might not show up, or what would concretely have to change for them to do so.

The design sprint tools did not help in counteracting this issue. Empathy maps and customer journeys, rather than grounding the problem in real lived experiences, simply reflected the company problem and its consequences.

The result was a problem that remained very broad for a large part of the sprint and therefore hard to tackle. Defined so generally, it made it challenging to identify potential solutions, leaving the problem feeling fuzzy and untouchable. If the teams had instead anchored the problem to a specific person and situation by asking what is holding someone back from participating, it might have been easier to identify potential solutions in the early stages.

4.2 Shallow Customer Understanding

Because the problem was understood only in broad terms, the understanding of the customer remained shallow too. As teams approached the challenge from a company-level perspective, the resulting target customer profiles were highly generic and failed to surface the concrete behaviours, frictions, or motivations that could inform a solution.

This was especially visible in the personas. Most represented a broad demographic group, such as a student or an elderly person, with general behaviours and loosely defined motivations. They offered a rough indication of who to target, but did not serve the core purpose of a persona for early-stage innovation: to uncover a concrete and meaningful problem from a specific person's perspective, sharp enough to guide aligned solution development.

The few teams that invested in more detailed personas did not necessarily do better. One team developed a well-constructed persona, which was an entrepreneur in a specific region, with a clear background, concrete motivations, even using a real person as its

basis. Yet when they explained the problem or began designing a solution, they reverted to treating it as a general challenge that potentially anyone could face. This suggests the problem itself was not specific or acute enough to stay anchored to that person, because it was too easy to generalise it.

As a result, across most of the teams, the persona provided some basic concreteness around who the target customer might be, but it did not push teams to think about concrete problems for concrete people. Since the final personas were largely kept in the background rather than placed at the centre of the work, they were seen more as a way of checking that the target customer could be interested, rather than as a tool for understanding the problem and defining a solution for them.

4.3 Absence of Meaningful Validation

A recurring weakness across teams was the absence of meaningful validation and true customer interactions. Most teams rarely spoke to real people and therefore relied entirely on their imagined persona as a reference point, treating assumptions about customer needs as established facts.

For the teams that did engage with customers, validation remained superficial. Rather than opening with open-ended questions to uncover the problems customers actually face, teams presented their proposed solution and asked whether customers would be interested in it. This is a common failure pattern in early-stage innovation: when a solution is validated before the underlying problem, the feedback customers provide is typically of limited value. Even if a person expresses interest in a solution, without an urgent need their willingness to adopt the product often remains low.

Another issue was overreliance on secondary research. Some teams relied more on online sources than on direct customer conversations, which gave them only a surface-level understanding of the problem domain and did not bring them close enough to specific customers to assess whether the problem was acute enough to be worth solving. Secondary research may confirm that a problem exists at a general level, but it cannot determine whether it is urgent enough for a particular person to act on.

A final pattern was the use of personal experience as a substitute for validation. In one team, a member's familiarity with a specific context served as an implicit stand-in for customer research. Using personal experience as the only way to validate a problem creates an illusion of deep understanding, even if it reflects only one perspective at one point in time and embeds the person's own biases. It is therefore rarely a reliable proxy for the needs of a broader customer group.

4.4 Solutions Disconnected from Real Needs

Because the problem was never made concrete, the customer never truly understood, and assumptions never validated, the solutions teams developed were oftentimes vague. Without a specific person and a specific problem to design for, there was no meaningful constraint on what a solution should do or for whom it should solve a problem, making solution development challenging.

This resulted in two distinct types of outcomes. Some teams developed solutions that were logistically or organisationally infeasible, proposing ideas whose complexity and cost had no relationship

Table 1: Industry Challenges Addressed in the Students Sprint

ID	Sector	Challenge Summary
C1	Tourism	Reducing seasonal demand imbalance by attracting visitors to a regional destination year-round.
C2	Manufacturing	Monetising high-value vehicle batteries during their months-long seasonal idle period.
C3	Financial Services	Making cooperative bank membership feel relevant and engaging to younger generations.

to what the company could realistically adopt. Others remained at a high level of abstraction, such that their solutions were, for instance, marketing strategies for attracting potential customers without addressing any concrete need. In both cases, the solution reflected the vagueness of everything that preceded it.

4.5 AI as a Comfort Blanket

We also observed several teams turning to AI to navigate the fuzziness of the early innovation stage, using it for brainstorming, domain research, or exploring the feasibility of their ideas. However, rather than helping them move toward a more concrete understanding of the problem, AI largely reflected their existing thinking back to them. When given a vague input, it returned a vague output with more words and a different formulation, elaborating on whatever assumptions the team already held without adding more content.

As a result, AI gave teams the feeling of progress and comfort without the substance of it. Fuzzy thoughts, once elaborated by an AI, can appear more legitimate and developed than they are, which may have reinforced teams' tendency to remain at the surface rather than pushing them to go deeper. As such, the problem was not in the tool per se, but in the expectation that it could resolve a conceptual challenge that the teams themselves had not yet worked through.

5 Design Implications

Based on the observations above, we derive two sets of design implications for a human-AI collaborative artefact intended to support novice teams in early-stage innovation.

5.1 Implications for Problem Framing

The findings point to a consistent problem across teams: the inability to achieve meaningful customer-problem fit before moving toward solution development. In general, early-stage innovation requires teams to remain in uncertainty long enough for a real, worthwhile problem to emerge. Most teams we observed, however, did not uncover a meaningful problem for a concrete customer before starting solution development. This was largely because they did not engage directly with customers or investigate them through other forms of research, which would have helped them concretize their understanding and surface a true problem more effectively.

The results also point towards the design tools not being able to help them find more concrete problems. When teams approached empathy maps or personas from the framing of the challenge, the tools simply reproduced that framing in a more structured format. Even if the technique was followed correctly, it did not allow them

to actively think about a concrete, meaningful problem for a targetable and realistic customer. As such, even by complying with the technique, it does not actively guide teams towards the right path.

This has a direct implication for what kind of support teams need: The core challenge already appeared in the earliest stage of innovation, namely understanding the problem and for whom it arises. An effective support artefact should therefore intervene at this stage, guiding teams toward uncovering a specific customer and a clearly defined, burning problem that is urgent enough to matter in that person's day-to-day life. For this, it should also help them understand how to conduct the necessary research, by helping in interviews or observational studies.

5.2 Implications for AI Collaboration

The findings also reveal the role of AI in early stage innovation. While LLMs were used by many teams, they never provided truly impactful help in their innovation processes.

This limitation emerged directly from how AI was used. Early-stage innovation demands critical questioning, such as asking whether a problem is real, whether the assumed customer actually exists, and whether the problem framing reflects genuine insight or simply mirrors inherited assumptions. These questions require teams to slow down and interrogate their assumptions before moving forward. Nevertheless, AI reflected back the same level of abstraction and substance they brought to it.

That said, there is still value in what AI does well, namely making thoughts explicit, capturing and structuring ideas, and elaborating on initial directions. For AI to be genuinely useful at this stage, however, it would need to go beyond elaboration and actively guide teams toward more concrete thinking, prompting them to move from abstract challenge framings toward specific customers and specific problems worth solving.

6 Artefact Concept

Based on these design implications, most teams appear to have failed at the *Understand* step of the design sprint, as they were unable to reach genuine customer-problem fit. The artefact we are aiming to develop in this research therefore initially targets this step, aiming to help teams navigate the earliest phase of innovation.

To support teams here, we reframe the *Understand* step through the lens of design thinking, where it corresponds to *empathize*². In this framing, the goal is to observe the problem from the perspective of specific people, their feelings, their context, and the concrete situations in which the problem arises, so that teams develop a

²<https://ixdf.org/literature/topics/empathize>

specific and grounded understanding of which problems are truly meaningful. With this, the problem becomes less fuzzy and more approachable by teams.

Empathy is typically built through direct engagement with customers, primarily via interviews and observation. This serves two purposes: it surfaces real information about the problem, and it forces teams to confront whether the problem they are solving actually exists in the form they imagined. Grounding the problem in real-world context makes it more concrete and opens up clearer paths toward potential solutions.

To guide teams in achieving customer–problem fit, we plan to develop a lightweight web application that uses storytelling as the primary method for building empathy. The application is intended to be used collectively by the full innovation team on a single shared screen, reflecting the collaborative setting observed in the field study. It is planned to guide teams through the following approach:

6.1 Uncover Assumptions through Stories

The first phase aims to surface the assumptions teams already hold about their customers and the problems they face. This is achieved by eliciting specific stories from teams that capture these assumptions. By telling realistic stories about concrete customers, their daily lives, and when and how the problem occurs, teams make the problem more tangible and counteract the fuzziness of early-stage innovation.

Because the story must remain plausible, it also serves as an early reality check, prompting teams to question whether their assumptions hold before speaking with real customers. It also implicitly checks for urgency, as a trivial problem tends to make the story feel flat or unconvincing.

AI plays an active role in this process by supporting narrative construction through interrogative prompts. By asking open-ended questions, the system encourages users to think through concrete story elements, including who the target customer is, what their daily context looks like, and where specific problems arise, thereby eliciting greater detail and specificity. After the story is elicited, the AI visualizes the underlying assumptions, making them explicit and testable. At this stage, teams may interact with the system by writing or speaking to reduce friction as much as possible.

6.2 Validation Through Primary Research

Once assumptions about the problem are made explicit, the team is in a position to test them against reality. Through primary research, particularly empathy interviews with potential customers or observational studies, teams can explore whether the story they constructed reflects reality or is pure fiction. Speaking to a real person who resembles the character they sketched gives the abstract narrative more depth and reveals where the assumptions were right, where they were incomplete, and where they were wrong entirely.

AI continues to play a collaborative role in this phase. It can coach teams in designing effective research studies, for instance by helping them design interviews or observational studies that generate genuine insights and clarifying what to watch for at each stage of the study. Where relevant, it can also draw on secondary

sources to find supporting data, adding a quantitative layer to the qualitative insights gathered through conversation.

6.3 Insight Elaboration

Once the primary research is completed, the team returns to the developed story and refines it with the newly gathered insights. As such, the assumptions that were validated are now explicitly confirmed and new insights are incorporated. Wrong assumptions are also made explicit, showing clearly where the initial ideas of teams were misleading.

The AI can accelerate this process by helping to analyse interview transcripts or observation notes, identifying patterns and surfacing relevant findings. It can also flag methodological concerns, for instance pointing out where an insight appears to stem from a leading question rather than a genuine response, and help the team assess how well-founded their understanding actually is.

Table 2 summarises how the observed challenges map to design principles and corresponding artefact mechanisms.

7 Evaluation Plan

Once the first proof-of-concept version is ready, we will evaluate it with novice innovation teams in an entrepreneurial university course. The tool will be used as part of their regular project work on real innovation projects, allowing us to assess it under realistic conditions.

We plan to conduct a longitudinal multiple-case study in which each team is observed using the tool over three months. During this time, we aim to iteratively evaluate and improve the artefact and deepen our understanding of the underlying problem, through the method described in the following subsections.

7.1 Data Collection

Data will be gathered through diverse methods. (1) Teams will be *observed* while using the tool, capturing how they interact with it, where they engage most, and where friction arises. (2) Following the session, *semi-structured interviews* will be conducted to explore their experience in more depth. Finally, (3) for usage outside the classroom, we will track usage behaviour to understand potential usage patterns and engagement over time. Tracked metrics will include session frequency, duration, and feature interactions, collected via anonymised logs within the tool itself.

7.2 Evaluation Dimensions

The collected data will be evaluated mostly focusing on three dimensions. The first is the concreteness of the customer definition that emerges from the session, specifically whether teams move from a broad target group to a specific person with a specific situation, which was the breakdown point identified in field observations. The second is the quality of problem understanding, and whether teams finish the session feeling closer to a real, grounded problem rather than a general challenge area. The third is the overall experience of the tool, including whether teams found it useful, whether the conversational approach felt natural, and whether the shift from conversation to visualisation supported their thinking.

The findings will be used continuously to assess what each artefact iteration gets right and where it falls short, directly informing

Table 2: Mapping observed challenges to design principles and artefact mechanisms.

Observed challenge	Design principle	Artefact implementation
High-level problem framing and shallow customer understanding	Force concreteness through realistic scenarios	“Day in the life” story elicitation
AI mirrors vague input (“comfort blanket”)	Use AI as an interrogative collaborator	Probing questions and specificity prompts
Weak or missing validation	Elicit assumptions and test them with real evidence	Primary research guidance

the next design and development cycle. They may also reveal new failure patterns and misunderstood challenges in early-stage innovation, guiding the next iteration of both the research project and the artefact.

8 Discussion

8.1 Relation to Design Thinking and Hypothesis-Driven Development

The proposed artefact builds directly on the empathize step of design thinking [2], which proposes first understanding the user’s lived experience, then defining a problem, and only later generating solutions. For the first iteration of the artefact, we therefore focus on developing a deep understanding of users and their problems, using a technique similar to “a day in the life” from user research³.

Furthermore, the structure of the artefact, which first surfaces assumptions, then tests them through primary research, and later refines understanding accordingly, maps closely onto hypothesis-driven approaches such as the Build-Measure-Learn loop [15] and related Lean Startup thinking. A key distinction, however, is that these frameworks are most commonly applied at the solution level, helping teams test whether a product works, while this artefact currently targets initial problem exploration before building solutions.

8.2 When to Build

Positioning the artefact before solution development raises the question of when building becomes appropriate. Traditionally, problem understanding was treated as a prerequisite, partly because development was expensive. However, as building becomes faster and cheaper through tools such as vibe coding, the temptation to skip problem understanding may grow stronger. Yet the argument for finding a meaningful problem before building likely still holds regardless of cost, since otherwise teams may simply build the wrong product faster. That said, cheaper building opens a new possibility: a prototype can also serve as an instrument for problem exploration, suggesting a future extension of the artefact that supports teams in exploring the problem through a solution.

9 Summary and Next Steps

In summary, this paper identifies a persistent challenge in early-stage innovation, especially for novice innovation teams: even with established methods being applied, teams frequently achieve only a superficial understanding of the problems they are addressing. It

also shows that current AI tools tend to reinforce rather than challenge the vague assumptions that underlie this. Building on these observations, we propose an initial prototype concept, supporting teams in surfacing, structuring, and validating their assumptions about customers and the problems they face.

The immediate next step is to implement and evaluate this prototype with novice innovation teams, assess whether it leads teams to more concrete customer and problem definitions, and improve the prototype design for future iterations of the study.

References

[1] Volker Bilgram and Felix Laarmann. 2023. Accelerating innovation with generative AI: AI-augmented digital prototyping and innovation methods. *IEEE engineering management review* 51, 2 (2023), 18–25.

[2] Tim Brown. 2008. Design Thinking. *Harvard Business Review* 86, 6 (2008), 84–92.

[3] Ahmed Fawzy, Amjed Tahir, and Kelly Blincoe. 2025. Vibe Coding in Practice: Motivations, Challenges, and a Future Outlook—a Grey Literature Review. *arXiv preprint arXiv:2510.00328* (2025).

[4] Carmine Giardino, Xiaofeng Wang, and Pekka Abrahamsson. 2014. Why early-stage software startups fail: a behavioral framework. In *International conference of software business*. Springer, 27–41.

[5] Alan R Hevner. 2007. A three cycle view of design science research. *Scandinavian journal of information systems* 19, 2 (2007), 4.

[6] Alan R Hevner, Salvatore T March, Jinsoo Park, and Sudha Ram. 2004. Design science in information systems research1. *MIS quarterly* 28, 1 (2004), 75–106.

[7] Anil Khurana and Stephen R Rosenthal. 1997. Integrating the fuzzy front end of new product development. *Sloan Management Review* 38, 2 (1997), 103–120.

[8] Jinho Kim and David Wilemon. 2002. Focusing the fuzzy front-end in new product development. *R&D Management* 32, 4 (2002), 269–279.

[9] Jake Knapp, John Zeratsky, and Braden Kowitz. 2016. *Sprint: How to solve big problems and test new ideas in just five days*. Simon and Schuster.

[10] Peter Koen, Greg Ajamian, Scott Boyce, Allen Clamen, Eden Fisher, Steve Fountoulakis, Albert Johnson, Pabjit Puri, and Robert Seibert. 2002. Fuzzy Front End: Effective Methods, Tools, and Techniques. In *The PDMA ToolBook 1 for New Product Development*, Paul Belliveau, Abbie Griffin, and Stephen Somermeyer (Eds.). John Wiley & Sons.

[11] Peter A Koen, George M Ajamian, Scott Burkart, Allen Clamen, Rachel Davidson, Ron D’Amore, Charlie Elkins, Kendra Herald, Michael Incorvia, Albert Johnson, et al. 2001. Providing clarity and a common language to the “fuzzy front end”. *Research-Technology Management* 44, 2 (2001), 46–55.

[12] Ash Maurya. 2012. *Running lean : iterate from plan A to a plan that works*. O’Reilly, Sebastopol, CA.

[13] Maha Mohamed Alsebai Mohamed, Pingfeng Liu, and Guihua Nie. 2022. Causality between technological innovation and economic growth: Evidence from the economies of developing countries. *Sustainability* 14, 6 (2022), 3586.

[14] Dan Olsen. 2015. *The lean product playbook: How to innovate with minimum viable products and rapid customer feedback*. John Wiley & Sons.

[15] Eric Ries. 2011. *The Lean Startup: How Today’s Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business, New York.

[16] Deborah L Roberts and Marina Candi. 2024. Artificial intelligence and innovation management: Charting the evolving landscape. *Technovation* 136 (2024), 103081.

[17] Ranjan Sapkota, Konstantinos I Roumeliotis, and Manoj Karkee. 2025. Vibe coding vs. agentic coding: Fundamentals and practical implications of agentic ai. *arXiv preprint arXiv:2505.19443* (2025).

³<https://think.design/user-design-research/a-day-in-the-life/>