



Fakultät für Ingenieurwesen
Facoltà di Ingegneria
Faculty of Engineering

Bachelor in Informatics and Management of Digital Business

Bachelor Thesis

How to reach Problem Solution Fit?

AI Powered Build-Measure-Learn loops

Candidate: Sebastiano Druda

Supervisor: Dr. Eduardo Martins Guerra

Co-Supervisor: Dr. Xiaofeng Wang

Bozen, July 2026

Abstract

The growing integration of artificial intelligence into entrepreneurial practices is reshaping how early-stage startups discover, validate, and develop new ideas. Understanding how AI can support the search for Problem Solution Fit is essential, as many startup failures originate from solving poorly framed or non-existent problems.

This thesis examines the challenge of reaching Problem Solution Fit by analysing how AI can enhance the Build-Measure-Learn loop during the earliest phases of a venture. The study adopts a mixed methodological approach that combines conceptual analysis, external case study examination, and exploratory first-hand case studies, alongside the evaluation of AI-assisted frameworks used in real startup environments.

The results show that AI can significantly accelerate hypothesis generation, data analysis, and early experimentation, helping founders validate problems more quickly and with greater precision. At the same time, the findings highlight the need for disciplined problem framing and human interpretation to avoid misleading signals amplified by AI.

For founders facing similar challenges, the thesis suggests that AI is most effective when used to support structured learning rather than to replace it. Startups that combine AI-driven insights with rigorous experimentation and direct user understanding are better positioned to reach Problem Solution Fit and to build solutions with scalable potential.

Riassunto

La crescente integrazione dell'intelligenza artificiale nelle pratiche imprenditoriali sta trasformando il modo in cui le startup nelle fasi iniziali scoprono, validano e sviluppano nuove idee. Comprendere come l'IA possa supportare la ricerca del Problem–Solution Fit è fondamentale, poiché molti fallimenti derivano dalla risoluzione di problemi mal definiti o inesistenti.

Questa tesi analizza la sfida di raggiungere il Problem–Solution Fit esaminando come l'intelligenza artificiale possa migliorare il ciclo Build–Measure–Learn nelle primissime fasi di un'iniziativa imprenditoriale. Lo studio adotta un approccio metodologico misto che combina analisi concettuale, l'esame di casi studio esterni e casi esplorativi di prima mano, insieme alla valutazione di framework supportati dall'IA utilizzati in contesti startup reali.

I risultati mostrano che l'IA può accelerare in modo significativo la generazione di ipotesi, l'analisi dei dati e la sperimentazione iniziale, consentendo ai founder di validare i problemi in modo più rapido e preciso. Allo stesso tempo, i risultati evidenziano la necessità di una rigorosa definizione del problema e di un'interpretazione umana attenta, per evitare segnali fuorvianti amplificati dall'uso dell'IA.

Per i founder che affrontano sfide simili, la tesi suggerisce che l'IA è più efficace quando viene utilizzata per supportare un apprendimento strutturato piuttosto che per sostituirlo. Le startup che combinano insight guidati dall'IA con una sperimentazione rigorosa e una comprensione diretta del comportamento degli utenti sono meglio posizionate per raggiungere il Problem–Solution Fit e sviluppare soluzioni con potenziale di scalabilità.

Zusammenfassung

Die zunehmende Integration von Künstlicher Intelligenz in unternehmerische Praktiken verändert grundlegend, wie Startups in frühen Phasen neue Ideen entdecken, validieren und weiterentwickeln. Das Verständnis darüber, wie KI die Suche nach dem Problem–Solution Fit unterstützen kann, ist von zentraler Bedeutung, da viele Startup-Misserfolge auf schlecht definierte oder nicht existente Probleme zurückzuführen sind.

Diese Arbeit untersucht die Herausforderung, Problem–Solution Fit zu erreichen, indem analysiert wird, wie KI den Build–Measure–Learn-Zyklus in den frühesten Phasen einer Unternehmung unterstützen kann. Die Studie verfolgt einen gemischten methodischen Ansatz, der konzeptionelle Analyse, die Untersuchung externer Fallstudien sowie explorative Fallstudien aus erster Hand kombiniert und diese mit der Analyse KI-gestützter Frameworks aus realen Startup-Kontexten verbindet.

Die Ergebnisse zeigen, dass KI die Generierung von Hypothesen, die Analyse von Daten und frühe Experimente erheblich beschleunigen kann, wodurch Gründer Probleme schneller und präziser validieren können. Gleichzeitig verdeutlichen die Ergebnisse die Notwendigkeit einer disziplinierten Problemdefinition und menschlicher Interpretation, um irreführende Signale zu vermeiden, die durch KI verstärkt werden können.

Für Gründer, die vor ähnlichen Herausforderungen stehen, zeigt diese Arbeit, dass KI am wirkungsvollsten ist, wenn sie zur Unterstützung strukturierten Lernens eingesetzt wird und dieses nicht ersetzt. Startups, die KI-gestützte Erkenntnisse mit rigoroser Experimentierung und direktem Verständnis des Nutzerverhaltens kombinieren, sind besser positioniert, um Problem–Solution Fit zu erreichen und skalierbare Lösungen zu entwickeln.

Acknowledgements

I would like to express my deepest gratitude to the people who have supported me throughout this journey.

My first and heartfelt thanks go to my beloved girlfriend, who has been by my side in every moment, including the most difficult ones. Her constant support, her encouragement to rediscover the value of studying as a path for personal growth, and her gentle insistence on helping me slow down my work pace so I could find a healthier balance have been truly invaluable.

I am sincerely grateful to my family for their emotional, financial, and motivational support during these years of study. A special mention goes to my dachshund, Pongo, whose visits always brought joy, lightness, and a welcome break from the demanding rhythm of work and academic responsibilities.

I also want to thank all the professors who generously shared their knowledge and guided my learning. Their dedication has played an essential role in shaping my academic experience.

My appreciation extends as well to my “desk companions”, who made this experience complete and memorable, turning daily study into a shared and enjoyable journey.

Finally, yet no less importantly, I wish to thank my supervisor Dr. Eduardo Martins Guerra for his professionalism and availability and my co-supervisor, Professor Xiaofeng Wang, for her guidance, her patience, and her support in the development of this thesis. Her advice has been fundamental in bringing this work to completion and in marking the conclusion of a path that has been deeply meaningful to me.

List of Figures

1.1	Overview of the thesis methodology	4
2.1	Problem-Solution Fit Canvas. Reproduced from: Problem-Solutions fit Canvas 2.0, IdeaHackers [1]	8
2.2	Build-Measure Learn Cycle. Reproduced from: The lean Startup, Eric Ries [2]	10
3.1	Product-Market Fit Pyramid. Reproduced from: The lean Product Play- book, Dan Olsen [3]	17
3.2	Double Diamond Diagram. Reproduced from: The Double Diamond, Design Council [4]	19

Contents

1	Introduction	1
1.1	Motivation: The Age of AI-Augmented Startups	1
1.2	Objectives and Research Questions	2
1.3	Structure and Methodology	2
1.3.1	Selection and documentation of literature sources	5
2	Background Information	7
2.1	What is a Startup	7
2.2	What is Problem-Solution-Fit	8
2.3	What is the Lean (Startup) Methodology	10
2.4	What are AI and LLMs	11
3	Practical Dimension	13
3.1	Secondary Case Studies on AI-Native Products	13
3.2	Why the Problem phase dominates the Product phase	16
3.3	AI as a learning accelerator in PSF frameworks	17
3.3.1	AI Enhances Problem-Solution-Fit Canvas and Boosts Lean Cycle	17
3.3.2	Starting the "Revamped Double Diamond framework" in the right way	18
3.4	Primary Case Study: Two AI-Supported Side Projects	21
3.4.1	FiscoBud: Problem discovery from direct experience	21
3.4.2	Bundyclock: Problem discovery inside an organisation	22
3.4.3	Cross-case insights	23
4	Discussion	25
4.1	Key Findings and Insights	25
4.1.1	Recommendation for SaaS Founders and Innovators	27
4.1.2	How to enhance outcomes through intelligent, faster, and mo- re concise lean cycles.	29
4.2	Personal Considerations	29
4.3	Reply to the Research Question	30
5	Conclusion and Further Studies	33
5.1	Final Summary	33

Chapter 1

Introduction

1.1 Motivation: The Age of AI-Augmented Startups

Today's startups grow in an environment where Artificial Intelligence is part of everyday work. Tools based on machine learning and Large Language Models help founders analyse markets, interact with potential users, and test ideas in ways that were much slower and more expensive in the past. These tools make it more trivial to generate assumptions, create early prototypes, and learn from feedback. As a result, new ventures can explore opportunities with a level of speed that was not even imaginable a few years ago.

At the same time, this rapid progress creates new challenges. AI can support teams in making sense of complex information, it can also generate misleading patterns or encourage founders to trust early inputs without checking by looking for real-world evidence. Many teams rely on the same AI tools, and because these models recombine patterns from existing data, the outputs tends to be similar to each other, making it harder for a startup to build something truly distinct.

As a result, researchers and practitioners have started to discuss these signals, pointing out that the benefits of AI come together with important risks that need to be understood, especially in the early steps of venture creation.

This tension between opportunity and uncertainty is the starting point of this thesis. The first stage of a startup, Problem-Solution Fit (referred to as PSF), is where many important choices are made. If teams misunderstand the problem or validate it poorly, later efforts often fail even if the product is well built. Understanding how AI can support or influence this phase is therefore essential. This thesis examines AI not only as a source of efficiency, but also as a force that shapes how founders think, learn, and make decisions in the earliest moments of a newly founded venture.

1.2 Objectives and Research Questions

The goal of this thesis is to explore how Artificial Intelligence can help startups reach Problem-Solution Fit by improving the Build-Measure-Learn loop used in the Lean Startup methodology. Much of the existing literature discusses Product-Market Fit or focuses on AI in later phases of growth. This thesis, instead, concentrates on the first stage, where teams define the problem, test their assumptions, and evaluate whether their proposed solution addresses a real need.

The core research question is the one proposed as the title of the thesis: **How to reach Problem-Solution Fit?**, and it is explored through two guiding sub-questions:

1. How do AI tools influence the way founders identify, understand, and validate the problem they want to solve?
2. How can AI improve early experiments and feedback cycles in a way that supports clearer and faster progress toward Problem-Solution Fit?

Together, these questions guide an examination of whether AI can reduce uncertainty in an early-stage venture, by enabling reliable problem framing and faster feedback cycles, ultimately improving the startup's ability to reach PSF.

1.3 Structure and Methodology

In order to examine how AI supports startups in reaching PSF, this thesis follows a structured path that progresses from contextual framing, to theoretical foundations, to practical investigation, and finally to discussion. This progression integrates two complementary components: a systematic review of relevant literature and an analysis of real case studies. Together, these perspectives allow for a balanced view that connects scientific evidence with practical application.

The first methodological component consists of a review of peer-reviewed articles, grey literature, industry research, and practitioner tools related to AI, Lean Startup practices, early validation, and Problem-Solution Fit. The selection of sources was guided by targeted keyword combinations, including “Problem-Solution Fit”, “AI accelerated discovery”, “Build-Measure-Learn”, “Lean Startup validation”, “LLM early stage experimentation”, “startup failure analysis”, and “AI for customer discovery”. These keywords were used across academic databases such as Google Scholar, SpringerLink, ScienceDirect, ACM Digital Library and other reputable industry platforms including Harvard Business Review, Startup Genome and Miro's market research library.

Furthermore, sources were selected according to two criteria. First, they needed to directly address the concepts central to this thesis, such as startup experimentation, problem framing, or AI-assisted validation. Second, they needed to represent

different viewpoints. This choice was deliberate: adopting diverse perspectives reduces the risk of bias and makes it possible to understand the topic from a variety of optimistic or critical academic, industrial and practitioner angles. This balance ensures a more complete understanding of the opportunities and limitations of AI.

The second methodological component consists of a set of case studies (referred to as "Secondary Case Studies") of, current or at the time, startups that used AI in their discovery and validation processes. These case studies were identified through literature scanning, targeted searches, and cross-referencing between articles. The selection process followed three steps. First, a research for real, documented cases where AI played a clear role in early validation, using keywords such as "AI startup case", "AI validation failure", "AI go-to-market example", and "AI powered discovery". Second, the cases have been filtered according to the relevance of their early-stage learning process, choosing those in which AI had a measurable, visible impact. Third, ensuring that the cases represented different outcomes: successful companies that achieved Problem-Solution Fit, companies that struggled despite strong technology and failed ventures that illustrate common pitfalls.

Each case study was analysed through the same lens: how the startup defined the problem, how it generated hypotheses, how it designed experiments, how it interpreted user feedback, and how it adjusted its strategy. This structure made it possible to extract comparable insights from different situations. The analysis focused on identifying patterns across successes and failures, clarifying what AI made easier, what it accelerated, and where it created misleading signals. These findings were then linked back to the frameworks discussed in the conceptual part of the thesis.

In addition to external, documented startup cases, the methodology was extended to include two personal side projects developed by the author: *FiscoBud* and *Bundyclock*. These projects (referred to as "Primary Case Studies") were analysed as reflective exploratory cases and were not intended to provide statistical generalisation. Instead, they serve as complementary empirical material that illustrates how AI-supported discovery is in practice when the founder is directly involved in problem identification and solution design. The same analytical lens was applied to these cases, focusing on problem framing, hypothesis generation, experimentation and learning dynamics. Their inclusion strengthens the methodology by connecting theoretical frameworks and external evidence with first-hand observation of AI-assisted Problem-Solution Fit processes.

By combining structured literature analysis with case-based investigation, the methodology provides conceptual grounding and documented evidence. The research path reflects the idea that understanding early validation requires attention to theory, to real examples, and to the interaction between human reasoning and AI-assisted processes.

The following diagram provides a visual overview of the methodological path adopted in this thesis. It summarises the main stages of the research process and shows how literature analysis and case study investigation were combined to answer the research question.




Figure 1.1: Overview of the thesis methodology

1.3.1 Selection and documentation of literature sources

To ensure transparency and methodological rigor, all literature sources used in this thesis were systematically documented and categorized according to their academic status, role in the research, and credibility. The literature base was divided into two main groups: peer-reviewed academic literature and grey literature.

Academic literature

As introduced before, the academic literature used in this thesis consists of peer-reviewed journal articles, academic books, selected working papers and institutional reports. Peer-reviewed articles form the theoretical and empirical backbone of the study, particularly in the areas of Lean Startup, entrepreneurial learning, startup failure and Artificial Intelligence. Academic books and papers are included as foundational references, given their central role in shaping widely adopted frameworks such as the Lean Startup methodology, Problem-Solution Fit, and Value Proposition Design.

In addition, a limited number of working papers and institutional publications are used to reflect recent developments in AI research and policy, where peer-reviewed evidence is still emerging. All sources were selected based on relevance, reliability and alignment with the research objectives.

Grey literature

Grey literature was included to capture practitioner knowledge and industry evidence that reflect real-world startup practices not yet fully covered by academic research. Given the exploratory nature of early-stage innovation and AI adoption, such sources provide valuable contextual insight into how Problem-Solution Fit is approached in practice.

To ensure transparency and methodological control, all grey literature sources were explicitly documented and evaluated using the following structured table. Each source was classified by type, publisher, publication date, role within the thesis, and credibility signals. In addition, the table clarifies why each source was considered relevant and how it contributes to the analysis. This approach makes the use of non-peer-reviewed material explicit and reduces the risk of implicit bias by clearly distinguishing empirical observation from theoretical validation.

Ref.	Source type	Author / Organisation	Year	Role in thesis	Reason for selection and discovery process
[4]	Framework reference page	Design Council	—	Theoretical and practical basis of the Double Diamond framework	Identified as the authoritative institutional source of the Double Diamond framework. Selected due to its recognised credibility and relevance to structured problem exploration and solution development.
[10]	Industry report (PDF)	Startup Genome	2021	Industry evidence on startup ecosystems and failure patterns	Identified through cross-referencing academic literature on startup failure. Selected for its large empirical base, structured methodology, and frequent use as an industry benchmark in both research and practice.
[13]	Practitioner article (blog)	Daria Nepriakhina	2017	Conceptual support on the Problem–Solution Fit Canvas	Identified through targeted keyword searches on Problem–Solution Fit frameworks. Selected for its practitioner-oriented interpretation of the PSF Canvas, which complements academic definitions with operational insights widely adopted in startup practice.
[21]	Case analysis (growth / marketing article)	Foundation Inc. Lab	2021–2022	Success case (Grammarly) illustrating PSF to PMF transition	Identified via targeted searches on AI-driven SaaS growth cases. Selected due to the lack of detailed academic case studies on Grammarly and its clear reconstruction of learning and validation loops.
[22]	Practitioner interview / article	Eilon Reshef (Gong)	2022	Success case (Gong) with focus on discovery discipline	Identified through venture capital and product discovery literature. Selected as a primary source authored by a company founder, offering direct insight into early problem framing and validation practices.
[23]	Trade press analysis	MM+M	2021	Failure and struggle case (Watson for Oncology)	Identified through searches on AI failures in regulated industries. Selected for its analytical focus on validation and implementation issues rather than technical performance alone.
[24]	News report	Nathan Rennolds	2023	Failure case (Zume) illustrating problem–solution mismatch	Identified via business journalism sources covering startup shutdowns. Selected for its clear timeline, identified author, and relevance to AI-driven business model misalignment.
[25]	Investigative tech journalism	Varsha Bansal	2025	Failure case (Builder.ai) and AI-washing pitfalls	Identified through longitudinal monitoring of AI startup narratives. Selected for its investigative depth and use of testimonials in a context where internal company data is not publicly accessible.
[26]	Investigative tech journalism	Grace Brown	2023	Failure case (Babylon Health) and warning on AI hype	Identified through investigative reporting on AI healthcare startups. Selected due to its depth of analysis and contextualisation of AI-related risks in sensitive domains.
[38]	Practitioner guide / vendor content	Miro	2025	Tooling example: AI for market research and team alignment	Identified during exploration of AI-supported discovery tools. Selected to illustrate how AI is operationalised in collaborative product discovery and research workflows.
[39]	Tool documentation	EdrawMax	2025	Example of AI-assisted canvases within the tooling ecosystem	Identified through comparative analysis of AI-enabled business design tools. Selected to represent emerging trends in AI-supported strategic and visual frameworks.

Table 1.1: Grey literature sources, bibliographic references, and selection rationale

Chapter 2

Background Information

The purpose of this chapter is to present the main ideas that form the foundation of this thesis and to explain why each of them is essential for understanding how Artificial Intelligence can support early-stage learning in startups. The next four sections introduce what a startup is, outline what AI and large language models can do, describe the Lean Startup methodology and its Build–Measure–Learn cycle and finally explain the concept of Problem-Solution Fit, which is the core focus of this research.

2.1 What is a Startup

Before exploring how AI affects early venture development, it is essential to clarify what is meant by the term “startup”. In academic literature, a startup is often defined not simply as a new business, but as a specific type of organization that operates under “high uncertainty” and focuses on testing and discovering a viable path to growth. One widely used definition states that a startup is a temporary organization created to search for a repeatable and scalable business model [5]. This description highlights that startups are built to learn, rather than to execute a fixed plan.

Other researchers take a more structural perspective. Luger and Koo describe a startup as a newly formed and independent business that is developing a viable product or service and is not the result of a merger, split, or expansion of an existing firm [6]. This definition focuses on the characteristics of the organization rather than its learning process, yet it still underlines the early and uncertain nature of the venture.

Practitioners also note an important distinction between startups and local small businesses. High-growth startups aim to build innovative solutions that can scale across markets, while local businesses typically pursue stable and geographically limited operations. Books such as *Startup Opportunities: Know When to Quit Your Day Job* [7] emphasize this difference by showing that a startup is defined not only

by being new, but by the ambition to grow quickly and serve a large market.

Taken together, these perspectives show that a startup is best understood as a young, independent, and learning-oriented organization that seeks to validate a scalable opportunity. This view helps explain why methodologies such as the Lean Startup and AI-powered tools are so relevant to early-stage teams: both are designed to support learning in environments where information is limited and decisions have long-term consequences.

2.2 What is Problem-Solution-Fit

Problem–Solution Fit describes the stage in a startup’s development where founders obtain convincing evidence that a meaningful customer problem exists and that their proposed solution has the potential to address it effectively. At this point, the startup is not yet concerned with market growth or product performance. Instead, the goal is to validate that the problem is real, relevant, worth solving and that the solution concept resonates with the targeted users. Maurya, in *Running Lean*, describes this moment as the point at which the team can confidently state that customers care about the problem and that the proposed approach can feasibly relieve the pain or meet the need [8].

Problem Solution Fit canvas

Customer segment(s)	Customer limitations	Available solutions Pros / Cons
Problems / Pains Its Frequency	Problem root / Cause	Behavior Its Intensity
Triggers to Act	Your solution	Channels Online
Emotions Before / After		Offline

Figure 2.1: Problem-Solution Fit Canvas.

Reproduced from: Problem-Solutions fit Canvas 2.0, IdeaHackers [1]

The concept has been further developed by Osterwalder and his colleagues in

Value Proposition Design. Their framework introduces the idea that the problem must be understood in terms of customer jobs, pains, and gains, and that the solution must provide clear pain relievers and gain creators. According to this perspective, Problem–Solution Fit occurs when these two maps overlap strongly: the customer’s underlying motivations and challenges align with the specific ways in which the solution promises to create value [9]. This alignment marks the transition from abstract idea to validated opportunity.

A growing body of evidence shows that this early stage is one of the most decisive in the entire startup journey. [10] Research often referenced within Harvard Business School and the widely cited *Startup Genome Report* both highlight that the majority of startup failures originate not from technical shortcomings but from inadequate validation of the problem itself. Many teams rush to build a product without confirming whether the problem truly matters to the users they hope to serve. [11] Harvard analyses, including the article *Why Startups Fail*, emphasize that early-stage teams frequently fall in love with their solution, assuming demand instead of demonstrating it. This leads to products that are well engineered but misaligned with actual customer needs.

Harvard Business Review reinforces this point in one influential article. In *Why the Lean Startup Changes Everything* [12], Steve Blank stresses that no methodology can compensate for a poorly defined problem. The Lean approach depends entirely on the assumption that the team is asking the right questions. Furthermore it argues that the first responsibility of a startup is to verify whether the problem is important before committing to development. If this validation is missing, even rapid experimentation risks producing misleading signals rather than useful learning.

Professional literature highlights these concerns. Sources such as [13] *Ten Insights of the Problem-Solution Fit Canvas and Pitfalls to Avoid When Seeking Problem-Solution Fit* reveal common mistakes: founders often misinterpret early customer enthusiasm as strong evidence, validate only with friendly audiences or assume that the problem is significant simply because the team finds it interesting. These materials underline that Problem-Solution Fit is not an intuition but a structured process of discovering why the problem exists, who experiences it and how urgently it needs to be solved.

More recent work combining The Lean Startup, Design Thinking, and the Jobs-To-Be-Done framework offers additional insights. [12] Haaverstad’s hybrid approach shows that a deep understanding of the problem reduces ambiguity in later development stages. When the problem definition is strong, experiments become easier to design, assumptions become clearer to test, and teams avoid drifting into product decisions based on speculation. In this sense, Problem-Solution Fit functions as an anchor that holds the entire innovation process together.

Because this thesis examines how AI can strengthen and accelerate the Build–Measure–Learn loop, the importance of Problem–Solution Fit becomes even more evident. AI can help founders generate hypotheses, analyse user behaviour, simulate scenarios and gather feedback rapidly. However, speed does not guarantee correctness. If the underlying problem is poorly defined or not meaningful, AI simply accelerates progress in the wrong direction. The quality of the learning process depends on the clarity and accuracy of the initial problem framing. This makes Problem–Solution Fit the essential foundation for understanding the role of AI in early-stage entrepreneurship.

2.3 What is the Lean (Startup) Methodology

The Lean Startup methodology provides a structured way for entrepreneurs to learn under uncertainty. Its core idea is that a startup succeeds when it can learn what works and what does not through continuous experiments. Ries defines the Lean Startup as a method for developing products and businesses that shortens development cycles and reveals whether a proposed business model is viable through validated learning, scientific experimentation, and iterative product releases [2].




Figure 2.2: Build-Measure Learn Cycle.

Reproduced from: The lean Startup, Eric Ries [2]

At the center of this approach is the Build-Measure-Learn cycle. According to Ries, the main activity of a startup is to transform ideas into products, measure how customers respond, and then learn whether to continue in the same direction or pivot toward a new one [2]. This cycle places learning above execution. Its purpose is not to perfect a solution on the first attempt, but to run small and frequent tests that produce evidence.

The Lean tradition has earlier roots in Lean Thinking, which promotes the idea of creating value while reducing waste. Womack and Jones describe Lean Thinking as a method that focuses on maximizing customer value with fewer resources [14]. Ries built on this principle by applying it to entrepreneurship and combining it with the Customer Development process introduced by Steve Blank, which stressed the importance of discovering customer needs before building a product.

Literature continues to confirm the relevance of the Lean Startup approach, noting that continuous experimentation helps teams reduce risk and make informed decisions even in complex environments. Research such as *How Do Entrepreneurs Think They Create Value?* [15] also shows that the Lean mindset encourages founders to adopt a scientific attitude toward learning and decision-making.

For these reasons, the Lean cycle is not simply a useful framework. It has become an essential practice for startups that want to understand their customers, test assumptions early, and adjust their strategy with evidence rather than intuition.

2.4 What are AI and LLMS

In recent years, the term Artificial Intelligence has become widely used and extensively studied, leading researchers to propose multiple definitions.

In general, Artificial Intelligence is a broad field that studies systems capable of perceiving information and acting upon it. Russell and Norvig describe AI as the study of agents that receive inputs from the environment and perform actions that influence it [16]. This definition captures the idea of AI as a set of techniques that help machines make sense of the world and produce useful outputs.

Large language models (LLMs) represent one of the most influential developments in this field. They are neural networks trained on massive collections of text to predict the next word or token in a sequence. Brown and colleagues describe them as self-predictive models built to learn linguistic and factual patterns from large datasets [17]. More recent work refers to them as foundation models, since they can adapt to a broad range of tasks once trained at scale [18].

In simpler words, according to the OECD, an AI system is a machine-based system that derives how to generate outputs such as predictions, recommendations, content, or decisions that can influence physical or digital environments [19]. This definition is particularly relevant because it focuses not only on how AI works, but also on the real-world impact it produces. In fact, beyond the strictly technical descriptions, AI and LLMs are technologies that are able to influence how people make decisions and handle personal or work-related tasks.

Clearly, their extensive use in every context and domain represents a great opportunity for increased productivity. For startups, specifically, AI can boost traditional product development methods and can reduce the costs of experimentation. Nevertheless, their consistent use implies responsibility and judgment.

In certain circumstances, if not enough nuance or crucial detail is provided when engineering a prompt, LLMs can provide biased or incorrect outputs and even “hallucinations”, which can distort facts and the reliability of their answers. Whether that is generated text or code, the user should take responsibility for validating and verifying outputs at all times in such a way to confidently ensure that a given response is accurate and can be trusted [20].

For example, in the case of producing code to implement a new software feature, especially in a startup context, validation becomes a two-sided requirement: it depends both on what the user can competently judge and what the model can accurately provide. First, careful validation requires sufficient knowledge on the domain and human judgment to confirm whether the code generated is accurate and aligned with the implementation of the feature. If the code provided by an LLM is outside of the user’s experience and domain, then validation is not entirely possible because they cannot reliably detect subtle errors, assumptions or missing compliance. On the model side, LLMs can still struggle in contexts characterized by uncertainty and incomplete information. When deductions and intuitive conclusions are needed, which can be the case when market needs and customer’s problems are being investigated for the deployment of an innovative feature, Artificial Intelligence provides limited and unreliable responses.

When users attribute too much competence to a model and when validation is weak, the risk is not only building a product with technical failures, but also a solution that misreads market needs, ultimately delaying the PSF and worsening the “high uncertainty” condition that characterizes the nature of startups.

Chapter 3

Practical Dimension

The practical dimension of this thesis examines how real startups have used Artificial Intelligence during the early stages of validation and why this matters for reaching Problem-Solution Fit. The intention of this chapter is not only to analyze what worked and what failed in these cases, but also to show how AI can strengthen the Problem phase, accelerate the Build-Measure-Learn cycle, and support structured innovation frameworks such as the Revamped Double Diamond. By connecting practical evidence with theory, this section demonstrates why a deep understanding of the problem, supported by reliable data and well designed experiments, is a necessary foundation for creating solutions that can scale.

3.1 Secondary Case Studies on AI-Native Products

Real case studies provide concrete evidence of how AI can help startups uncover real customer problems, refine their solutions, and avoid mistakes that often lead to failure. For this thesis, the case studies were selected through a structured methodological approach based on keyword-driven searches as explained in the Structure and Methodology chapter. Each of them has been selected in order to have a credible analyses to ensure accuracy and a wide range of perspectives. This approach made it possible to build a set of cases that differ in outcome, industry, and level of success, allowing a fair and unbiased comparison of how AI affects discovery and Problem-Solution Fit. By comparing both successful and unsuccessful applications, it becomes possible to extract detailed lessons that can guide founders who want to use AI responsibly and effectively in the early stages of their venture.

A strong example of a successful application is Grammarly, which built a multi-billion dollar business by addressing a clear and high-frequency writing problem. The company focused on grammar correction, clarity suggestions, tone adaptation, and writing style improvement. Grammarly's early strategy combined machine learning with continuous observation of real user behaviour. Because writing errors occur frequently and are easy to recognise, the problem did not require complex discovery work. The feedback loop was immediate: users could see the value

of suggestions instantly, which created a reliable validation mechanism during onboarding. Grammarly also used anonymised writing patterns to refine its models continuously, making AI training a direct extension of customer understanding. The company prioritised measurable outcomes such as error reduction, clarity improvement, and rewriting speed, which made it possible to verify Problem-Solution Fit early. This case reveals that when the problem is concrete and measurable, AI can accelerate learning by reducing uncertainty and giving teams fast and reliable evidence [21].

Another positive example is Gong, a company that leveraged AI to analyse real sales conversations at scale. Gong used speech recognition and natural language processing to detect patterns in customer behaviour, such as recurring objections, common requests, effective messages, and emotional responses. Instead of relying on isolated feedback from sales representatives, Gong built a structured database of actual conversations. This generated insights that teams could use to prioritise features, refine the value proposition, and improve messaging. Gong's approach demonstrates how AI can turn unstructured human communication into actionable customer signals. The company also adopted a disciplined discovery process: before committing to new features, it examined whether patterns found by AI were confirmed by direct interaction with design partners and customer interviews. This prevented teams from acting on misleading correlations. Gong's success is therefore not only technical, but also organisational, showing how AI and validated learning can reinforce one another [22].

Not all applications of AI succeed. IBM's Watson for Oncology [23] is one of the most well-known cases of an AI-driven initiative that struggled to reach Problem-Solution Fit. Watson was built to provide clinical recommendations for cancer treatment, but its development relied heavily on limited training datasets and assumptions that did not generalise to real medical environments. Many of its recommendations were inconsistent with local clinical practices or not aligned with individual patient histories. Moreover, clinicians often found that the system did not integrate naturally into their workflow, creating friction and reducing trust. The technology was advanced, but the team lacked a deep understanding of the variability of clinical decision-making and the real needs of doctors. As a result, Watson often produced advice that clinicians rejected, revealing a fundamental misalignment between the solution and the real problem context. This case shows that AI cannot compensate for weak problem discovery, especially in high-stakes domains where accuracy, context sensitivity, and user trust are essential.

Several other examples strengthen similar insights. Zume Pizza [24], a highly funded venture that attempted to reinvent the pizza industry with robots and automated delivery systems, collapsed despite its technological sophistication. The company focused on automation efficiency rather than validating whether customers appreciated robot-made pizza or whether such a model was economically sus-

tainable. This illustrates how AI and automation can become distractions when teams neglect to validate the core problem. Builder.ai [25], another interesting example, struggled because it relied heavily on synthetic assumptions about customer needs. The company used AI to generate app prototypes quickly, but failed to confirm whether users actually wanted fully automated app creation or whether these prototypes matched real business objectives. Babylon Health [26], an AI-driven healthcare platform, faced similar issues; it relied on algorithmic predictions that were not always aligned with real patient behaviour or clinical constraints. These examples show that AI can create misleading confidence when it is used to extract needs from data without direct validation through interviews, field experiments, or real user observation. Research on entrepreneurial failure further confirms that ventures often collapse due to misreading the problem or assuming demand too early, rather than due to technical limitations [27].

The literature provides further evidence by showing how AI can strengthen early discovery and opportunity identification. Existing analysis demonstrates that founders use AI tools to scan large datasets, interpret early signals, and generate preliminary hypotheses more quickly than traditional manual research. However, the same research emphasises that AI improves discovery only when insights are validated through human interpretation and real-world engagement [28].

Quantitative data from the CB Insights “Top Reasons Startups Fail” dataset further reinforces this point. According to this widely cited report, the top cause of failure is “no market need”, which directly reflects a lack of validated Problem-Solution Fit. Teams often build products based on assumptions rather than evidence, leading to solutions that do not solve a meaningful problem. This dataset offers numerical evidence consistent with the failed cases reviewed in this thesis [29].

Harvard Business Review provides an additional perspective on how AI should be integrated into early-stage discovery work. Recent analyses argue that human-AI collaboration produces better strategic decisions than either component alone, especially in uncertain and rapidly changing environments. AI speeds up analysis, while human founders contribute contextual understanding and domain intuition. This helps explain why Grammarly and Gong succeeded while Watson for Oncology struggled: in the successful cases, AI insights were always interpreted and validated by humans, while in the failed cases, teams relied too heavily on model-driven assumptions without sufficient real-world evidence [30].

A deeper comparison of all cases reveals several detailed patterns. Teams that succeeded treated AI as an assistant rather than a decision-maker. They used AI to speed up data processing, identify early signals, and explore hypotheses, but they always verified insights with customers or design partners. Successful cases also treated AI outputs as suggestions rather than facts, which reduced the risk of acting on misleading correlations. Moreover, these teams paired AI insights with metrics

that reflected real user behaviour, such as engagement, retention, conversion, or observed behavioural change. On the other hand, the failed cases tended to rely too heavily on AI-generated insights, especially when data were synthetic, incomplete, or unrepresentative. They often skipped basic discovery steps such as deep interviews, problem mapping, or field observation, assuming that AI could provide accurate guidance without the need for human interpretation.

From these cases and additional sources, a clear conclusion emerges. AI can accelerate go-to-market learning by analysing large volumes of data, identifying hidden patterns, and generating insights faster than manual methods. However, fast insights are only valuable when they match real customer needs. Teams that reach Problem-Solution Fit combine AI-driven analysis with direct validation and structured learning. Teams that fail often rely on AI as a substitute for discovery, which leads to inflated assumptions and solutions that do not match the real problem. The practical takeaway is clear: AI is a powerful assistant, but only when the problem is defined with care and validated through real evidence. When used responsibly, AI strengthens the Build-Measure Learn loop and helps founders reach Problem-Solution Fit faster and with more confidence. When used without discipline, AI accelerates the wrong type of learning and amplifies errors.

3.2 Why the Problem phase dominates the Product phase

The case studies discussed above show that many failures originate from mistakes in the Problem phase rather than from poor product execution. This aligns with research that emphasises the central importance of defining and validating the problem before building or scaling a solution. Studies on Lean Startup experimentation highlight that building quickly, without a clear understanding of the problem, often produces waste and false signals. According to Lean Startup and the Business Model: Experimentation Revisited, teams that rush into building features without strong problem framing risk misinterpreting customer behaviour and drawing conclusions from incomplete or misleading data [31].

Research on experiments in innovation further argues that well designed experiments reveal the structure of the problem itself, not only the performance of the product. Experimenting for Novelty and Impact shows that experimentation helps teams uncover assumptions that shape the problem, identify boundary conditions for their solution, and refine their understanding of customer needs [32].

Peer reviewed studies on startup failure also confirm that weak problem discovery is a core deficit. The article Why Do Startups Fail? A Core Competency Deficit Model identifies poor understanding of customer needs and market context as primary contributors to failure. The authors argue that without a validated problem, even highly capable teams with strong technical skills struggle to find traction [33].




Figure 3.1: Product-Market Fit Pyramid.

Reproduced from: The lean Product Playbook, Dan Olsen [3]

This perspective is reinforced by incubation research. The paper Proposed Method for Problem-Solution Fit Phase at Start Up Incubator shows that structured problem validation processes, including qualitative interviews, early testing, and iterative refinement, significantly increase the likelihood of success [34].

These findings make a clear case: the Problem phase is not only important, it is the foundation on which all later development depends. A strong Problem-Solution Fit ensures that resources are directed toward meaningful needs, experiments evaluate the right questions, and solutions are built on real insights rather than assumptions. For this reason, this thesis emphasises that founders should devote more attention to defining the problem than to building the product in the early stages.

3.3 AI as a learning accelerator in PSF frameworks

3.3.1 AI Enhances Problem-Solution-Fit Canvas and Boosts Lean Cycle

Artificial Intelligence provides a new layer of capabilities that can strengthen the Problem-Solution Fit Canvas and accelerate the Lean cycle. The Problem-Solution Fit Canvas is traditionally used to map customer pains, desired gains, key assumptions, and early solution ideas. AI can enrich each stage of this process by helping founders access broader data, generate hypotheses more quickly, and run experiments at lower cost.

One of the most practical applications is the use of AI for market research. A Har-

vard Business School working paper shows that LLMs can analyse thousands of customer interviews, reviews, and support tickets to extract recurring themes and jobs to be done [35]. This capability allows teams to identify patterns that would take weeks or months to collect manually. In addition, AI tools can summarise these insights into structured formats that feed directly into the Problem-Solution Fit Canvas.

Other studies, such as Predictions Through Lean Startup? Harnessing AI Based Predictions Under Uncertainty, explore how AI can generate and prioritise hypotheses. By ranking hypotheses according to frequency, severity, and potential business impact, AI helps founders decide what to test first. This aligns with research showing that experiments should focus on risky assumptions rather than on convenient or interesting ones [36].

AI also supports the Build-Measure Learn cycle by enabling faster measurement. For example, AI powered landing page generators, automated message testing, and predictive models can run micro experiments that provide early signals of interest. The article Artificial Intelligence, Lean Startup Method, and Product Innovations describes how AI can complement the Lean approach by producing more frequent learning cycles and reducing the cost of iteration [37].

Market research platforms, such as Miro's AI research guides, also show how companies use AI to accelerate time to Product-Market Fit by aligning teams on insights and shortening decision cycles [38]. The Lean Canvas ecosystem now includes AI assisted tools that help founders articulate assumptions, define customer segments, and simulate personas [39].

Taken together, these developments show that AI can improve the entire Problem-Solution Fit process: discovering pains, generating hypotheses, testing assumptions, and documenting insights. However, the value of AI depends on human interpretation. AI can suggest directions, but founders must confirm these signals through experiments, interviews, and real customer behaviour.

3.3.2 Starting the "Revamped Double Diamond framework" in the right way

The Double Diamond is a well known design process that divides innovation into four sequential phases: Discover, Define, Develop, and Deliver. Its purpose is to separate problem exploration from solution development, ensuring that teams do not rush into building products based on untested assumptions. The Revamped Double Diamond adapts this structure to contemporary digital innovation. It maintains the same foundation but places stronger emphasis on the early discovery of the problem before any solution is imagined or tested. This mirrors the distinction established earlier in this thesis between the Problem phase and the Product phase.

In both cases, meaningful innovation begins with an accurate and validated understanding of the problem. Without this foundation, even well-designed products risk solving issues that are not relevant in the reality.




Figure 3.2: Double Diamond Diagram.

Reproduced from: The Double Diamond, Design Council [4]

The first stage of the Revamped Double Diamond is entirely focused on understanding the problem space. In this phase, teams collect information, analyse user behaviour, engage with stakeholders, and investigate the underlying causes that give rise to the perceived problem. The goal is not to define a solution but to clarify what the real problem is and why it exists. Research from the Design Council confirms that the quality of the early discovery phase strongly influences the quality of later creative work, prototyping, and decision making. When teams begin the Double Diamond with a strong problem frame, they reduce ambiguity in downstream phases, simplify experimentation, and increase the likelihood of developing an effective and scalable solution [4].

Artificial Intelligence can significantly strengthen this early stage. Recent practical guides and playbooks show that AI can support founders and product teams throughout the discovery phase by amplifying their ability to analyse, classify, and synthesise information at scale [3]. The first step is to establish a clear mission that defines what the team wants to understand and why. This is followed by explicit validation criteria that determine what counts as evidence of a real problem. Once the mission and scope are defined, AI driven techniques can be applied to speed up the actual discovery process. Large language models and NLP tools can examine interviews, reviews, support tickets, social media posts, and recorded sales calls. These tools can summarise the most common pains, identify emotional triggers, detect behavioural patterns, and highlight emerging themes. This allows teams to explore a larger problem landscape than what purely manual research can typically uncover.

One of the most valuable contributions of AI in this phase is the creation of a Problem Framing document. This document synthesises the evidence gathered during discovery and provides a structured description of the problem, the people affected, the circumstances in which the problem emerges, and the implications of not solving it. AI can help draft early versions of this document by organising insights and extracting the most relevant information from large datasets. However, the final interpretation remains a human responsibility, ensuring that the resulting problem frame reflects real user behaviour rather than model generated assumptions.

Beyond analysis, AI plays a meaningful role in hypothesis generation. Large language models can propose a wide range of explanatory hypotheses about why the problem exists, what its root causes might be, and which contexts influence user behaviour. They can also generate early solution ideas that teams can use for feasibility discussions. These suggestions are not meant to replace creativity but to expand the initial set of ideas that teams can explore. AI can also simulate conversations with personas to help refine value propositions, anticipate objections, and improve early messaging before running costly or time consuming live tests.

Several tactical guides for applying AI to the Revamped Double Diamond outline clear and practical steps for founders. First, teams should focus their discovery process by collecting data from multiple sources and ensuring that these inputs reflect real user behaviour. Second, they should use AI to generate hypotheses and rank them based on potential impact, frequency, severity, feasibility, and expected willingness to pay. Third, teams should run small scale, low cost experiments such as landing page tests, A B tests, or short interviews to verify whether the hypotheses hold. Finally, they should converge toward a refined problem definition that reflects both human insights and AI generated evidence.

Research from the Hasso Plattner Institute and other design research institutions reinforces these practices. Studies show that when AI enhances discovery, the Double Diamond process becomes faster and more precise, but the benefits materialise only when human interpretation remains central. AI is highly effective at identifying patterns and exploring large volumes of data, but human judgment is responsible for validating these insights, avoiding biases, and making strategic decisions. This hybrid approach ensures that problem discovery benefits from computational speed while remaining grounded in real user needs.

Integrating AI also helps avoid a common error of traditional discovery work. In many cases, teams enter the Double Diamond with predefined ideas of the solution. This creates a risk of confirmation bias, where teams look for information that supports their assumptions rather than challenging them. AI mitigates this risk by expanding the search space and highlighting patterns that might not align with ini-

tial expectations. This pushes teams to reconsider their assumptions early, before they commit time and resources to the wrong direction.

Once the problem is clearly framed, the Define phase becomes significantly more efficient. A validated problem frame guides the criteria for evaluating early concepts, selecting which ideas deserve deeper exploration, and designing appropriate experiments. AI can support this phase by analysing user segments, identifying similar cases, and offering different perspectives on the problem. As a result, the Define phase becomes more structured and less sensitive to subjective decisions driven by intuition alone.

Starting the Revamped Double Diamond in the right way therefore means grounding the entire innovation process in a validated and evidence based understanding of the problem. AI strengthens this foundation by expanding the founder's ability to explore, analyse, and synthesise information. It helps teams avoid premature solution development, reduce uncertainty, and create a clearer transition between problem discovery and solution creation. A stronger Problem phase increases the chances of identifying the right problem, designing a scalable solution, and avoiding expensive pivots later in the process.

When applied with discipline and supported by human interpretation, AI becomes a powerful partner in the Double Diamond framework. It accelerates discovery, strengthens validation, and reduces ambiguity, ultimately allowing founders to begin the innovation journey in a more informed and reliable way.

3.4 Primary Case Study: Two AI-Supported Side Projects

To complement the external case studies presented earlier in this chapter, this section introduces a primary case study based on two personal side projects developed in parallel: *FiscoBud* and *Bundyclock*. The purpose is not to present these projects as proven market successes, but to use them as practical evidence of how AI can influence the path toward Problem–Solution Fit (PSF) through faster Build–Measure–Learn loops and through a different validation logic depending on whether the problem emerges from direct experience (often described as “scratch your own itch”) or from observation inside an organisation. This perspective is aligned with the thesis focus on how AI supports early-stage learning rather than replacing it.

3.4.1 FiscoBud: Problem discovery from direct experience

FiscoBud originated from a direct operational need: managing both the economic (cash flow visibility, cost planning, renewals) and fiscal (documents, deadlines, compliance steps) aspects of an Italian *Partita IVA forfettaria* without delegating routine activities to a *commercialista*. The initial discovery step was not a generic idea-generation exercise, but a problem definition grounded in day-to-day friction: un-

certainty about which documents must be produced, in what order, and with what timing, combined with poor visibility of the financial consequences of those obligations.

The early PSF work followed three practical steps:

1. **Mapping the operational workflow.** The process was formalised as a sequence of events and documents (for example, quotation, acceptance, payment terms, invoice issuance, payment registration, tax-related tracking). This reduced ambiguity by transforming a vague need (“manage taxes”) into a set of explicit operational tasks that could be measured and improved.
2. **Competitor and substitute scanning.** Existing tools were evaluated primarily as partial substitutes, namely solutions that covered only one side of the workflow (either fiscal assistance or economic and administrative management), often designed for different target users (for example, larger companies, or single-purpose services). This comparison strengthened the PSF framing by clarifying the gap: not that no solutions exist, but that solutions exist while they fragment the workflow and do not match the needs of a *forfettario* user trying to operate autonomously.
3. **AI-supported specification and architectural decisions.** AI tools were used as accelerators for the translation from workflow to system design. Rather than using AI to invent the problem, AI was applied to structure requirements into a coherent specification, explore feasible architectures and technology choices that support incremental development, and speed up market research by identifying relevant players and categories. This produced a practical benefit aligned with the thesis findings: AI reduced the time needed to generate and organise hypotheses and implementation options, but it did not remove the need for human judgment in deciding what to build and why.

From a PSF perspective, the key limitation of this path is also clear. Internal validation (“it works for us”) is strong evidence that a problem exists for at least one realistic user, but it is not sufficient evidence that the same problem is urgent, frequent, and worth paying for across a broader segment. Early user interest (for example, sign-ups) may indicate curiosity, but it does not automatically demonstrate willingness to pay or sustained adoption. For this reason, FiscoBud represents a PSF path where AI mainly strengthens problem formalisation and iteration speed, while the most uncertain part remains external validation under real constraints such as trust, compliance concerns, and data security expectations.

3.4.2 Bundyclock: Problem discovery inside an organisation

Bundyclock originated from a different trigger: the observation of inefficiencies in an existing time-tracking system used in a real organisational context. The problem

was not defined as “build a time tracker”, but as reducing recurring operational friction caused by usability problems that generate mistakes (missed entries, inconsistent logs) and administrative overhead created by employee uncertainty (frequent requests to verify hours, corrections, and clarifications).

Here, the PSF reasoning was driven by workflow pain rather than market whitespace. The proposed value proposition focused on error reduction and clarity, implemented through design constraints that directly target the observed failure modes (for example, simplifying the entry and exit logic into a single interaction pattern to reduce inconsistent states). Additional requirements emerged from real work modalities (hybrid, on-site, remote), motivating features such as location-related verification when relevant and clearer weekly and monthly visibility for employees.

AI contributed differently in this second case:

- **Rapid synthesis of alternatives and feature benchmarking.** AI was used to accelerate the exploration of what existing tools typically provide, how competitors frame pricing models, and which features are commonly missing or implemented poorly. This supported faster hypothesis formation (for example, “simplicity and transparency reduce admin workload”) and faster iteration on scope.
- **Faster prototyping and debugging cycles.** AI-assisted development reduced build time and enabled a quicker build–measure rhythm during implementation. However, the core PSF signal still came from contextual evidence: the presence of concrete administrative costs (time spent on corrections, repeated employee questions) and the feasibility of adopting a simpler workflow.

Compared to FiscoBud, Bundyclock illustrates a PSF path where the initial evidence is more contextual and operational (the problem is observed in a real organisation), while the scalability question becomes more complex. A solution that works much better in one company may still be hard to adopt elsewhere, especially when existing systems are already in place and seen as “good enough.” This reinforces a central theme of the thesis: AI can accelerate iteration and research, but it cannot remove resistance to adoption or replace direct access to decision makers and their constraints.

3.4.3 Cross-case insights

Taken together, the two projects highlight a consistent pattern with the external cases discussed earlier: AI amplifies the learning system that founders already have. When the problem is grounded in concrete workflows, AI is highly effective at speeding up documentation, hypothesis generation, competitive scanning, and the construction of minimum artefacts that allow measurement. At the same time, both

projects expose the same boundary condition. PSF is not achieved by building faster, but by learning faster from signals that are connected to real behaviour, including continued use, reduced friction, willingness to pay, and trust in sensitive contexts.

In practical terms, this primary case study supports the thesis argument that AI is most valuable in the earliest phases when used to structure and accelerate discovery work, turning messy ideas into testable hypotheses and coherent experiments, while keeping the validation aspects based on real users, real constraints, and measurable outcomes.

Chapter 4

Discussion

4.1 Key Findings and Insights

The analysis developed throughout the practical dimension, combined with the theoretical background, provides a detailed view of how Artificial Intelligence can strengthen early-stage entrepreneurial learning, particularly during the pursuit of Problem-Solution Fit. This pattern is also reflected in the primary case studies presented in the Practical Dimension, where AI was primarily used to structure problem definitions, organise requirements, and accelerate early research activities, rather than to generate the initial problem itself. One of the central conclusions that emerges is that AI does not replace the Build-Measure-Learn loop. Instead, it changes the scale, speed, and nature of the signals that can be detected at each step. The core engine of Lean Startup remains the same, but AI enables founders to observe the problem space in ways that would otherwise be too time consuming, too complex, or too broad to explore manually. This shifts the loop from a slow and linear process into a more dynamic and continuous one.

The case studies examined in the practical dimension, including both secondary case studies and the primary case studies based on the author's side projects, provide valuable insight into how AI interacts with the learning process. Across all examples, from Grammarly to Gong, and from Watson for Oncology to companies such as Zume Pizza or Babylon Health, as well as the author's projects, a consistent pattern emerges: AI amplifies whatever learning system it is placed in. When founders approach discovery with clarity, discipline, and a strong grasp of the problem space, AI accelerates progress and reduces uncertainty. When founders begin with vague assumptions, weak problem framing, or excessive confidence in early ideas, AI accelerates confusion instead. This amplification effect is visible in almost every case examined.

Three broad insights emerge from this analysis.

The first insight is that AI is most effective when the problem is concrete, fre-

quent, and directly observable in user behaviour. Problems of this nature tend to have clear indicators of success, which makes the feedback loop more reliable. Grammarly illustrates this dynamic clearly. It targeted a universal and easily recognisable pain: unclear writing, grammatical mistakes, and communication inefficiency. Users did not need explanation or onboarding to recognise whether a suggestion improved clarity. This meant that AI could immediately validate whether the model produced value. The quality of the underlying problem shaped the quality of the feedback loop. Similarly, Gong worked within a domain filled with rich behavioural signals: recorded sales conversations. Every call contained expressions of user intent, objections, concerns, and moments of resonance. AI enabled the team to analyse these conversations at scale, but the foundation for this analysis was the transparency of the problem itself. In both cases, AI accelerated learning because the context allowed users to naturally reveal whether the solution was helpful or not. Clear problems generate clear feedback.

The second insight is that AI makes weak assumptions more dangerous. When teams rely on incomplete data, superficial understanding, or untested beliefs about user needs, AI does not correct them. Instead, it amplifies them by speeding up decision making, increasing the volume of interpreted signals, and creating a false sense of confidence. The Watson for Oncology case illustrates this risk. The system was developed with a limited set of training data and insufficient validation in real-world clinical contexts. Although the model produced recommendations quickly, many of these recommendations were misaligned with clinical reality. Doctors found that the system did not integrate into their workflow and often disagreed with its suggestions. The issue was not the lack of technical sophistication but the lack of a validated problem definition. The same pattern appears in Zume Pizza, Builder.ai, and Babylon Health. In all three cases, AI-driven systems produced outputs that seemed impressive on the surface but were not anchored in a real understanding of user behaviour. AI provided speed, but speed without direction resulted in costly misalignment. A similar limitation emerged in the primary case studies, where internal validation and early interest were insufficient to confirm Problem-Solution Fit, reinforcing that AI cannot compensate for untested assumptions or substitute direct exposure to real adoption constraints.

The third insight is that the frameworks examined, including the Lean Startup cycle, the Problem-Solution Fit Canvas, and the Revamped Double Diamond, all converge on the same central message: most of the meaningful learning happens during the Problem phase. AI strengthens this stage by increasing the founder's capacity to observe users, analyse large datasets, and explore hypothesis spaces that would be difficult to map manually. It does not make the exploration less necessary; it makes it more thorough. AI expands the scope of what can be discovered, but the quality of the resulting insights still depends on careful experiment design, disciplined iteration, and thoughtful interpretation by humans. AI does not replace the intuition, empathy, or contextual understanding that founders bring to the dis-

covery process. It complements them.

Taken together, these findings show that AI and the Build-Measure-Learn loop prove most effective when used:

- to clarify the structure of the problem before investing in complex features,
- to convert unstructured user feedback into behavioural patterns that guide prioritisation,
- to enable many small, inexpensive experiments that reduce uncertainty incrementally,
- to accelerate the synthesis of insights and make each iteration more informative.

The combination is less effective when:

- AI outputs are treated as definitive conclusions rather than as hypotheses,
- experiments rely on superficial metrics such as clicks or impressions,
- teams reduce direct contact with users and rely heavily on AI-generated signals,
- the domain requires deep contextual validation, such as health or finance, but the solution is deployed with limited field testing.

In summary, AI improves learning when founders use it to engage more deeply with users and their behaviour. It undermines learning when it replaces this engagement. This distinction was also evident in the two personal projects discussed earlier, where AI-enabled speed in building and iteration did not automatically translate into faster validation, further underscoring the difference between execution velocity and learning quality.

4.1.1 Recommendation for SaaS Founders and Innovators

Based on the observation of case studies and frameworks, several recommendations emerge for SaaS founders and innovators who want to integrate AI into their discovery process.

First, founders should begin with the problem, not with the model. Many failures start with an excessive fascination with what a model can do, rather than with a proved understanding of what users need. Durable ventures emerge when teams develop a deep connection with the user's pain and let that insight guide their design choices. AI should be introduced later to expand this understanding, not to

substitute it.

Second, founders should choose problems that are easy to recognise and evaluate. Problems that users can identify instantly reduce the cost of validation. When the problem is too abstract, too rare, or too dependent on complex interpretation, early testing becomes slower and more ambiguous. AI can help interpret complexity, but it cannot fix ambiguity caused by weak discovery.

Third, founders should instrument their real data sources as early as possible. This does not require mature products. Even limited prototypes, early interviews, sales conversations, and customer support interactions generate signals that can be analysed by AI. Real behaviour is always more informative than synthetic assumptions or theoretical models.

Fourth, AI should be treated as an hypothesis generator rather than a decision maker. LLMs are excellent at identifying possible explanations, user segments, or solution concepts, but these remain hypotheses until validated. The Build-Measure-Learn loop begins after AI, not with it. AI expands the hypothesis space, but the team decides what to test and how.

Fifth, founders should design experiments that test behaviour rather than intention. Many failures come from believing what people say rather than observing what they do. AI can help summarise interviews and extract themes, but behavioural experiments provide clearer and more reliable evidence. For SaaS products, this can include sign-ups, feature usage, pre-orders, or willingness to pay.

Sixth, founders should avoid misleading metrics amplified by AI. Models can generate large numbers of impressions, clicks, or positive responses, but these do not necessarily indicate Problem-Solution Fit. Founders must define the specific signals that reflect meaningful progress, such as retention, repeated use of a pain-relieving feature, or real willingness to pay.

Finally, founders should document their learning explicitly. AI can help organise information, but the interpretation of what has been learned remains the responsibility of the team. Keeping a structured record through tools such as the Problem-Solution Fit Canvas reduces the risk of repeating mistakes and ensures that learning compounds over time.

Taken together, these recommendations show that AI should be used as a multiplier of strong discovery practices, not as a shortcut that bypasses them. The most successful SaaS founders will likely be those who combine careful problem framing with intelligent use of AI tools.

4.1.2 How to enhance outcomes through intelligent, faster, and more concise lean cycles.

The Build-Measure-Learn cycle is the engine of the Lean Startup methodology, designed to reduce uncertainty and support continuous improvement. One of the central questions of this thesis is how AI can make this cycle faster, more efficient, and more insightful.

In the Build phase, AI reduces the cost of creating prototypes and artefacts. Early versions of landing pages, user messages, mockups, onboarding flows, and even simple interfaces can be generated in a matter of seconds. Instead of building full products, teams can produce lightweight simulations of their ideas and test core assumptions without heavy investment. The focus should always be on building the minimum artefact necessary to measure behaviour linked to the problem.

In the Measure phase, AI offers even more powerful contributions. Models can analyse conversion funnels, detect recurring issues, summarise qualitative feedback at scale, and cluster user interactions into distinct themes. These capabilities allow founders to connect patterns across different channels, such as support requests, customer interviews, or navigation behaviour. To keep the cycle concise, founders should avoid tracking too many indicators. Instead, they should identify a small set of metrics directly connected to the specific problem they are testing.

In the Learn phase, AI helps organise information by generating structured summaries, highlighting contradictions, and uncovering relationships that might not be visible at first glance. However, the final synthesis must be carried out by humans. Teams must interpret the evidence, update their problem framing, and decide whether to persevere, pivot, or reframe the problem. A useful pattern is to let AI summarise data and let humans interpret meaning.

When combined with discipline and clarity, AI enhanced Lean cycles become faster and more precise. Speed is no longer a risk but an advantage, because each iteration is grounded in validated learning rather than speculation. The entire process becomes a way to converge toward a scalable solution while continuously refining the understanding of the problem.

4.2 Personal Considerations

Working on this thesis has significantly deepened my understanding of startups, innovation processes, and the role of Artificial Intelligence in early-stage learning. At the beginning, it was tempting to view AI mainly as a way to gain speed or as an advantage that came from adopting cutting-edge tools. As I progressed through the research, I discovered that speed alone is not a benefit. Many of the failures ex-

amined in the case studies were not slow. They moved quickly, raised substantial investments, and deployed complex technology. Their failures came from a lack of validated understanding, not from a lack of ambition or execution capacity.

This realization reshaped my perspective on what early-stage innovation requires. I learned that the most fundamental work happens in the Problem phase, long before a product reaches the market. AI can help amplify this work, but it cannot replace the human responsibility of careful observation, empathy, and interpretation. This made me more aware of the limits of AI and of the importance of human intuition in entrepreneurialism.

I also became aware that the frameworks explored, such as the Lean Startup cycle, the Problem-Solution Fit Canvas, and the Revamped Double Diamond, are not rigid formulas. They are mental habits that encourage continuous learning. These frameworks remind us that innovation is not a one-time bet but a sequence of discoveries. AI fits naturally into this view because it expands our ability to learn, but it becomes dangerous when it encourages founders to treat assumptions as facts.

Finally, this thesis made me more conscious of the responsibility that comes with designing AI-driven solutions, particularly in sensitive contexts such as health-care or finance. The consequences of incorrect assumptions can extend beyond financial losses. For future work, this reminds me that enthusiasm for innovation should always be balanced with respect for context, ethics, and long-term impact.

4.3 Reply to the Research Question

The research question that guided this thesis was divided into two parts:

1. How do AI tools influence the way founders identify, understand, and validate the problem they want to solve?
2. How can AI improve early experiments and feedback cycles in a way that supports clearer and faster progress toward Problem-Solution Fit?

The findings from the theoretical background, the practical case studies, and the methodological analysis allow for a clear and structured answer to both components.

1. How AI influences problem identification, understanding, and validation

The evidence collected throughout this thesis shows that AI significantly expands the range of observations that founders can make during the Problem phase. Tools such as large language models, natural language processing systems, and machine learning classifiers enable teams to analyze interviews, reviews, sales conversations,

support tickets, and behavioral traces with much precision than manual methods allow. This supports three major shifts in early-stage discovery.

First, AI increases the amount of data that can be processed during problem exploration. Instead of relying on a limited number of interviews or case-based insights, founders can extract patterns from large datasets. This helps them see recurring pains, contextual triggers, and user motivations that would otherwise remain hidden.

Second, AI strengthens the clarity of problem understanding by helping founders cluster similar observations, detect inconsistencies, and identify underlying themes. This reduces the risk of making product decisions on noisy or biased signals. In cases such as Gong and Grammarly, AI contributed directly to a clearer definition of the problem by revealing what users repeatedly struggled with.

Third, AI accelerates problem validation by providing early indicators of whether the identified pain is real, relevant, and frequent. By analysing behavioural data at scale, AI helps teams confirm whether the problem appears consistently across different segments and circumstances. However, the analysis conducted also shows that AI can mislead founders when the underlying assumptions are weak or when data are unrepresentative, as in the Watson and Babylon Health cases. For AI to be effective, it must operate within a disciplined discovery process grounded in human interpretation and direct user interaction.

Overall, AI influences problem identification and validation by expanding founders' observational capacity, improving pattern recognition, and accelerating the discovery of meaningful insights. However, it does not remove the need for human judgment, contextual interpretation, and behavioral verification.

2. How AI improves early experiments and feedback cycles

The second part of the research question concerns the Build-Measure-Learn loop and the role of AI in accelerating validated learning. The cases and frameworks analyzed reveal several important contributions.

In the Build phase, AI reduces the effort required to create prototypes, landing pages, messages, user flows, and simple simulations. This allows teams to test assumptions quickly and inexpensively, focusing on the smallest artifact that enables behavioral measurement. AI-generated variations support A/B testing, message validation, and early framing experiments.

In the Measure phase, AI provides powerful analytical capabilities. It can segment behavioral data, summarize qualitative feedback, detect recurring issues in user journeys, and uncover correlations between actions and outcomes. This short-

ens the time needed to understand whether an experiment produced meaningful insights. Instead of waiting for manual analysis or interpreting isolated feedback, founders can rely on broader and more structured evidence.

In the Learn phase, AI helps to synthesize all available data by generating summaries, highlighting contradictions, and suggesting alternative interpretations. Teams can use these outputs to refine their hypotheses, update their problem framing, and decide whether to pivot or persevere. The speed of analysis does not replace critical thinking but enables more informed and timely decisions.

Across these phases, AI improves the quality and speed of learning by enabling many small, low-cost experiments rather than a few large ones. It helps founders detect weak signals earlier, validate hypotheses more rigorously, and avoid investing in solutions that do not address real problems. However, as demonstrated in failed cases, AI becomes counterproductive when teams rely on its outputs without validating them through user behavior.

Final synthesis

Taken together, the two parts of the research question lead to a coherent conclusion: startups reach Problem-Solution Fit with AI when they use these tools to strengthen, rather than replace, the core learning activities of early-stage entrepreneurship. AI widens the observation window in the Problem phase and accelerates insights in the Build-Measure-Learn cycle, but the quality of the results depends on disciplined problem framing, clear hypotheses, and continuous contact with users.

The path to Problem-Solution Fit is therefore not based on automation alone. It emerges from the combination of AI-enabled analysis and human-driven interpretation. When this balance is achieved, founders can discover real problems sooner, design more relevant experiments, avoid common failure patterns, and move toward a validated solution with greater clarity and confidence.

Chapter 5

Conclusion and Further Studies

5.1 Final Summary

This thesis set out to understand how Artificial Intelligence can support early stage startups in reaching Problem-Solution Fit through faster and more focused Build-Measure-Learn loops. The work followed a progressive path, beginning with the broader context of AI augmented entrepreneurship, moving through the conceptual foundations that shape early innovation processes, and finally examining practical evidence drawn from case studies, frameworks, and contemporary applications of AI.

The Introduction established the motivation for the research by situating the study within the current landscape of AI accelerated venture creation. Startups today operate in a context where powerful models can generate ideas, analyse data, and build prototypes at speeds that were previously impossible. This raised a central question: does this acceleration help founders reach Problem-Solution Fit more effectively, or does it risk amplifying poor assumptions? The Introduction also clarified the methodology adopted in this thesis, combining peer reviewed literature, grey literature, and case based insights to build a balanced understanding of how AI interacts with the early stages of venture development.

The Background Information section built the conceptual structure necessary to address this question. It clarified what defines a startup, distinguishing it from conventional businesses by its search for a scalable and repeatable model. It then explored Artificial Intelligence and large language models in accessible terms, framing them not only as technological tools but also as cognitive extensions that influence the way founders reason and make decisions. The section also explained the Lean Startup methodology, emphasising the logic of iterative learning embedded in the Build-Measure-Learn cycle. Finally, it introduced the central concept of Problem-Solution Fit, positioning it as the foundation on which meaningful Product-market Fit is later built. This created the theoretical base that guided the rest of the work.

The Practical Dimension translated these concepts into concrete analysis. Through real, documented, and first-hand examples, the study showed how AI can strengthen or weaken early stage learning depending on how it is used. In particular, the inclusion of the primary exploratory case studies allowed the research to observe how AI-supported discovery unfolds when founders are directly involved in problem framing, experimentation, and iteration. Successful cases, such as Grammarly and Gong, demonstrated that AI can support founders in discovering patterns, validating assumptions, and refining their solutions when the underlying problem is clear and measurable. Unsuccessful cases, including Watson for Oncology and other AI driven ventures, highlighted the risks of relying on technology without validating the problem in real contexts. The Practical Dimension also examined frameworks such as the Problem-Solution Fit Canvas, the Lean Cycle, and the Revamped Double Diamond. Together, these tools showed how AI can support each stage of early validation, from problem discovery and hypothesis generation to experiment design and synthesis of learning. Throughout this analysis, a consistent pattern emerged: AI enhances the learning process when it is used to support clarity, focus, and disciplined experimentation.

The Discussion section brought together these findings and connected them with personal reflections, insights from exploratory case studies, and practical recommendations. It emphasised that AI is not a shortcut and does not replace the need for careful problem framing, observation of real user behaviour, and thoughtful experiment design. The key insights suggested that the real power of AI lies in helping founders process more information, test more assumptions, and interpret more signals within each iteration of the Build-Measure-Learn loop. The Discussion also presented a direct answer to the research question, concluding that Problem-Solution Fit can indeed be reached more effectively with AI, provided that the technology is used to amplify the founder's learning process rather than automate it.

Looking at the overall path of the research, the thesis shows a clear progression from context to theory, from theory to practice, and from practice to reflection. Each section supports the next, forming a coherent study of how AI reshapes early stage entrepreneurship. The conclusion is that AI's contribution to early validation is significant but conditional, as confirmed both by documented startup cases and by first-hand exploratory evidence. When founders pair AI with disciplined learning, strong problem discovery, and a clear understanding of user behaviour, they can move faster and with greater precision. When these elements are missing, AI may accelerate the wrong type of progress.

The study therefore reinforces an essential message for innovators: the future of entrepreneurship belongs not to those who move the fastest, but to those who learn the fastest. Artificial Intelligence can help accelerate this learning, but only when founders remain deeply engaged with the problem they are trying to solve and the evidence that reveals whether their assumptions are true.

Bibliography

- [1] IdeaHackers Network. Problem-solution fit canvas, 2019.
- [2] Eric Ries. *The Lean Startup: How Today's Entrepreneurs Use Continuous Innovation to Create Radically Successful Businesses*. Crown Business, 2011.
- [3] Dan Olsen. *The Lean Product Playbook: How to Innovate with Minimum Viable Products and Rapid Customer Feedback*. John Wiley & Sons, 2015.
- [4] Design Council. The double diamond. Last Accessed: 11/05/2026. Available at: <https://www.designcouncil.org.uk/our-resources/the-double-diamond/>.
- [5] Steve Blank. *The Four Steps to the Epiphany: Successful Strategies for Products That Win*. John Wiley & Sons, 2020.
- [6] Michael I. Luger and Jun Koo. Defining and tracking business start-ups. *Small Business Economics*, 24(1):17–28, 2005.
- [7] Sean Wise and Brad Feld. *Startup Opportunities: Know When to Quit Your Day Job*. John Wiley & Sons, 2017.
- [8] Ash Maurya. *Running Lean: Iterate from Plan A to a Plan That Works*. O'Reilly Media, 2022.
- [9] Alexander Osterwalder, Yves Pigneur, Gregory Bernarda, and Alan Smith. *Value Proposition Design: How to Create Products and Services Customers Want*. John Wiley & Sons, 2015.
- [10] Startup Genome. The global startup ecosystem report 2021. Last Accessed: 11/05/2026. Available at: <https://startupgenome.com/report/gser2021>.
- [11] Tom Eisenmann. *Why Startups Fail: A New Roadmap for Entrepreneurial Success*. Crown Business, 2021.
- [12] Steve Blank et al. Why the lean start-up changes everything. *Harvard business review*, 91(5):63–72, 2013.
- [13] Daria Nepriakhina. 10 insights of the problem-solution fit canvas. Last Accessed: 11/05/2026. Available at: <https://medium.com/lean-startup-circle/10-insights-of-problem-solution-fit-canvas-6999f75154fb>.

- [14] James P. Womack and Daniel T. Jones. Lean thinking: banish waste and create wealth in your corporation. *Journal of the Operational Research Society*, 48(11):1148, 1997.
- [15] Dennis Lyth Frederiksen and Alexander Brem. How do entrepreneurs think they create value? a scientific reflection of eric ries' lean startup approach. *International Entrepreneurship and Management Journal*, 13(1):169–189, 2017.
- [16] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Pearson, 4 edition, 2021.
- [17] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020.
- [18] Rishi Bommasani et al. On the opportunities and risks of foundation models. Technical report, Stanford University, 2021. arXiv preprint arXiv:2108.07258.
- [19] Recommendation of the council on artificial intelligence. Technical report, Organisation for Economic Co-operation and Development (OECD), 2019. Policy recommendation.
- [20] Simon Thorne. Understanding the interplay between trust, reliability, and human factors in the age of generative ai. *International Journal of Simulation-Systems, Science & Technology*, 25(1), 2024.
- [21] Foundation Inc. Lab. How grammarly built a \$13b+ brand targeting simple typos. Last Accessed: 11/05/2026. Available at: <https://foundationinc.co/lab/how-grammarly-built-13-billion-dollar-brand/>.
- [22] Eilon Reshef. Gong's path to product-market fit — why this \$7b company still works with design partners. Last Accessed: 17/05/2026. Available at: <https://review.firstround.com/gongs-path-to-product-market-fit/>.
- [23] MM+M. Where watson went wrong. Last Accessed: 17/05/2026. Available at: <https://www.mmm-online.com/home/channel/features/where-watson-went-wrong/>.
- [24] Nathan Rennolds. The \$500 million robot pizza startup you never heard of has shut down, report says. Last Accessed: 17/05/2026. Available at: <https://www.businessinsider.com/robot-pizza-startup-zume-shutting-down-raised-500-million-softban-2023-6>.
- [25] Varsha Bansal. Inside the collapse of builder.ai: Was it even an ai company? Last Accessed: 18/05/2026. Available at: <https://restofworld.org/2025/builderai-ai-apps-downfall/>.

- [26] Grace Brown. The fall of babylon is a warning for ai unicorns. Last Accessed: 18/05/2026. Available at: <https://www.wired.com/story/babylon-health-warning-ai-unicorns/>.
- [27] Wolfgang Lattacher and Malgorzata Anna Wdowiak. Entrepreneurial learning from failure: A systematic review. *International Journal of Entrepreneurial Behavior & Research*, 26(5):1093–1131, 2020.
- [28] Yaser Hasan Al-Mamary. The transformative power of artificial intelligence in entrepreneurship: exploring ai’s capabilities for the success of entrepreneurial ventures. *Future Business Journal*, 11(1):104, 2025.
- [29] CB Insights. The top reasons startups fail. 2023. Industry report.
- [30] H. James Wilson and Paul R. Daugherty. Collaborative intelligence: Humans and ai are joining forces. *Harvard Business Review*, 2021.
- [31] Teppo Felin, Alfonso Gambardella, Scott Stern, and Todd Zenger. Lean startup and the business model: Experimentation revisited. *Long Range Planning*, 53(4):101889, 2020.
- [32] Nancy Bocken and Yuliya Snihur. Lean startup and the business model: Experimenting for novelty and impact. *Long Range Planning*, 53(4):101953, 2020.
- [33] Edit Szathmári, Zoltán Varga, Attila Molnár, Gergely Németh, Zsolt Péter Szabó, and Orhideea Edith Kiss. Why do startups fail? a core competency deficit model. *Frontiers in Psychology*, 15:1299135, 2024.
- [34] Adhi Antyanto Naratama and Nila Armelia Windasari. Proposed method for problem-solution fit phase at start-up incubator. *European Journal of Business and Management Research*, 4(2), 2019.
- [35] Julian Brand, Ayelet Israeli, and Daniel Ngwe. Using llms for market research. Technical report, Harvard Business School, 2023. Working Paper No. 23-062, SSRN.
- [36] Santo Raneri, Fabian Lecron, Julie Hermans, and François Fouss. Predictions through lean startup? harnessing ai-based predictions under uncertainty. *International Journal of Entrepreneurial Behavior & Research*, 29(4):886–912, 2023.
- [37] Xiaoning Wang and Lynn Wu. Artificial intelligence, lean startup method, and product innovations. *Management Science*, 2025.
- [38] Miro. Ai in market research: How it streamlines planning and aligns teams on insights. Last Accessed: 18/05/2026. Available at: <https://miro.com/ai/ai-market-research/>.
- [39] EdrawMax. Ai lean canvas generator. Last Accessed: 18/05/2026. Available at: <https://www.edrawmax.com/app/ai-lean-canvas/>.