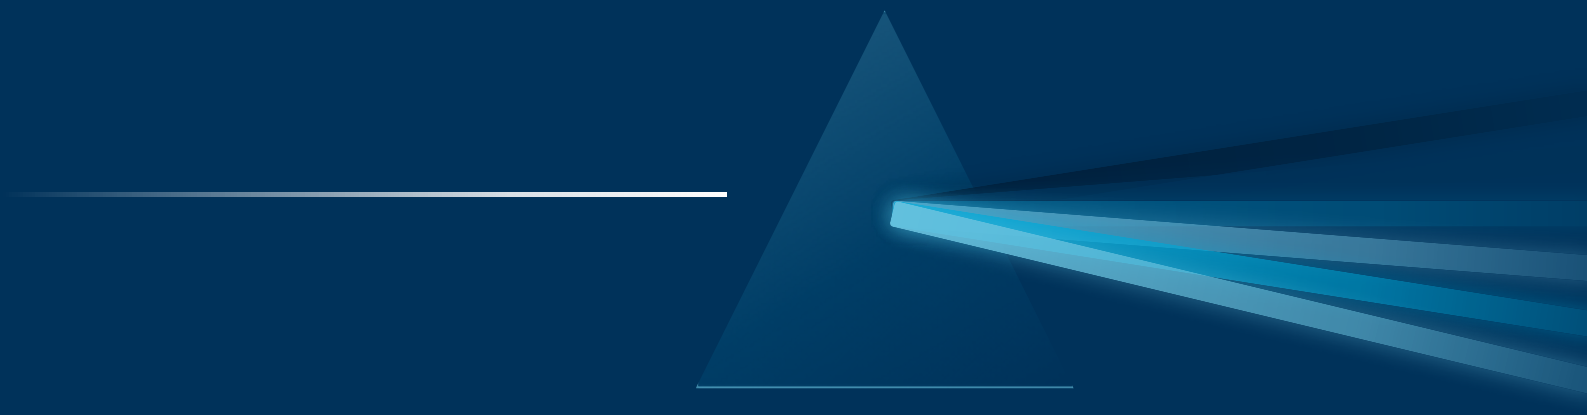
MSC GRADUATION THESIS

STRATEGIC PRODUCT DESIGN · TU DELFT



An AI-Enabled Sensemaking Framework for Go-to-Market decision makers.

Reducing information overload and improving insight quality for Product Marketing Managers at Webfleet

AUTHOR

Niranchna Balakrishnan

FACULTY

Industrial Design Engineering, TU Delft

IN COLLABORATION WITH

Webfleet — a Bridgestone company, Amsterdam Zuid

DATE

June 2026

i

An MSc graduation project at the Faculty of Industrial Design Engineering, TU Delft.

GRADUATION COMMITTEE

Graduation committee

This report documents an MSc graduation project carried out at the Faculty of Industrial Design Engineering, Delft University of Technology, in collaboration with Webfleet — a Bridgestone company in Zuid Amsterdam.

AUTHOR	**Niranchna Balakrishnan**
PROGRAMME	**MSc Strategic Product Design, TU Delft**
CHAIR	**Dr. Shahrokh Nikou**
MENTOR	**Dr. Nikoloas Kyriakopoulos**
COMPANY MENTOR	**Zilvia Hsu — Webfleet PMM Team Lead**
DATE	**June 2026**

"The real work was deciding what mattered, what it meant, and how to say it clearly."

Preface

This thesis began with a problem I went looking for within the PMM (Product Marketing) team at Webfleet. In conversations with people who do Go-to-Market work, the difficulty was rarely a shortage of information. It was the opposite — there was always more market intelligence than anyone could reasonably hold in their head, and the real work was deciding what mattered, what it meant, and how to say it clearly to someone else. That gap between having information and making sense of it is what this project set out to address.

Working with Webfleet gave the question a concrete home. Rather than treating AI as a general productivity promise, I was able to see how Product Marketing Managers, Regional Market Managers, and the sales team actually moved from fragmented signals to decisions, and where that process breaks down. The PRISM(E) framework grew directly out of those conversations, not the other way around.

I am grateful to my chair, Dr. Shahrokh Nikou, and my mentor, Dr. Nikoloas Kyriakopoulos, for the steady guidance and for pushing me to keep the work both academically rigorous and practically useful. A special thanks to my chair for agreeing to join my thesis project at the last minute, I am truly grateful and of course to my mentor, who helped me find my chair. To my company mentor Zilvia Hsu and the wider PMM team at Webfleet, thank you for your time, your candour in the interviews, and your willingness to build alongside me in the co-creation sessions.

This research sets out to meet an immediate need within Webfleet's PMM team and, in doing so, to lay the groundwork for the company's future AI initiatives.

Niranchna Balakrishnan

Delft, June 3rd 2026

iii
KEYWORDS

sensemaking
information overload
generative AI
prompting
product marketing
Go-to-Market

ABSTRACT

Abstract

Product Marketing Managers (PMMs) in B2B technology companies sit at the intersection of many information streams.

At Webfleet, a Bridgestone company operating in European commercial-vehicle telematics, PMMs must turn fragmented inputs from market research, customer feedback, sales signals, competitor intelligence, and regional sources into coherent Go-to-Market (GtM) insight. The difficulty is not a lack of information but an excess of it, and current practice relies heavily on individual judgment that is hard to scale, audit, or improve.

This thesis investigates how a structured, AI-enabled sensemaking framework can reduce information overload and improve the quality of GtM insight synthesis for PMMs at Webfleet. Following a qualitative, design-led approach grounded in design science research, the project ran across three phases: a contextual inquiry using semi-structured interviews with PMMs and secondary stakeholders, the design of a framework and AI prompt guide, and a perception-based assessment and trial use of the guide.

The central design outcome is **PRISM(E)**, a six-element framework: Prioritise, Reduce, Interpret, Synthesise, Mobilise, and an Equip layer of organisational enablement that runs alongside the five sequential stages. PRISM(E) is operationalised inside Microsoft 365 Copilot within Webfleet's own tenant — a choice grounded in data governance and infrastructure fit rather than raw capability. Interview findings refined the framework from its conceptual form into a final version, and the work is delivered through three artefacts: an AI Prompt Guide, a strategic three-horizon roadmap, and a tactical roadmap for adoption.

The contribution is a design-research-grounded operationalisation of sensemaking and prompt-engineering theory in the specific context of B2B product marketing — an area that existing literature leaves largely unaddressed.

iv

Terms used throughout this report, in the precise sense the framework gives them.

GLOSSARY

Glossary

PRISM(E)
The framework designed in this project — an acronym for Prioritise, Reduce, Interpret, Synthesise, Mobilise, with Equip as a sixth, non-sequential enablement layer.

Prioritise
The stage that filters and ranks incoming information against a stated GtM focus before any processing begins.

Reduce
The stage that compresses inputs to the few signals that genuinely require action, anchored to a specific decision rather than generic summarisation.

Interpret
The stage where compressed information becomes a structured insight draft, generated by AI under a B2B-analyst prompt.

Synthesise
The stage where the PMM adds context and judgment, runs an audit gate, and turns the draft into a shareable insight.

Mobilise
The stage that adapts the approved insight for different channels and archives it in a searchable knowledge base.

Equip
The organisational enablement layer addressing AI literacy, confidence, trust calibration, and prompt fluency.

GtM insight
An insight that is actionable, evidence-based, and explains the why behind a signal, formatted for a specific audience.

Audit gate
A five-question checkpoint within Synthesise that verifies tool fit, data handling, source attribution, human review, and governance comfort before an insight is shared.

Co-creation
The participatory design activity through which practitioners shaped and refined the framework and prompt library.

v

Acronyms and research vocabulary; this is the key.

ABBREVIATIONS

Abbreviations

AI	Artificial Intelligence
B2B	Business-to-Business
DSR	Design Science Research
GDPR	General Data Protection Regulation
GtM	Go-to-Market
LLM	Large Language Model
M365	Microsoft 365
PMM	Product Marketing Manager
RMM	Regional Market Manager
RQ	Research Question
SQ	Sub-Question
TAM	Technology Acceptance Model
TOE	Technology–Organisation–Environment
xAI	Explainable Artificial Intelligence

CONTENTS

Table of contents

Graduation committee 2

Preface 3

Abstract 4

Glossary 5

Abbreviations 6

1 Introduction 8
- 1.1 Webfleet, a Bridgestone company 9
- 1.2 Problem context 9
- 1.3 Research question and sub-questions 10
- 1.4 Competitive landscape and market context 11

2 Literature Review 12
- 2.1 Information overload in organisational contexts 12
- 2.2 Sensemaking in organisations 13
- 2.3 Prompt engineering and generative AI 14
- 2.4 AI in innovation management and B2B knowledge creation 15
- 2.5 Situated AI, Explainable AI (xAI), and digital transformation 16

3 The PRISM(E) Framework 18
- 3.1 Developing the framework 19
- 3.2 P — Prioritise 20
- 3.3 R — Reduce 20
- 3.4 I — Interpret 21
- 3.5 S — Synthesise 21
- 3.6 M — Mobilise 22
- 3.7 E — Equip 22

4 Methodology 24
- 4.1 Research approach 25
- 4.2 Contextual inquiry (Phase 1) 25
- 4.3 Concept development (Phase 2) 26
- 4.4 Evaluation (Phase 3) 26
- 4.5 Sampling strategy 27
- 4.6 Ethical considerations and data governance 27

5 Findings from the Contextual Inquiry 28
- 5.1 Analysing the interview data 29
- 5.2 Respondent characteristics 30
- 5.3 The stakeholder sample 31
- 5.4 How overload is experienced 31
- 5.5 A shared standard for a good insight 32
- 5.6 Trust in AI 32
- 5.7 Data treatment 33
- 5.8 Structural friction beneath the information problem 33
- 5.9 A shared demand for prompting support 34
- 5.10 Four gaps and the design targets 34

6 Refined Framework and Design 36
- 6.1 From design targets to a refined framework 36
- 6.2 Stage-by-stage refinements 37
- 6.3 Designing the prompt library 38
- 6.4 Co-creation with practitioners 39

7 Validation of the Prompt Guide 40
- 7.1 Validation approach and rationale 41
- 7.2 Primary validation: Product Marketing Managers 42
- 7.3 Secondary validation: Regional Market Managers 43
- 7.4 Secondary validation: Sales 44
- 7.5 Cross-role validation synthesis 45
- 7.6 Limitations and next validation steps 45

8 Final Deliverables 46
- 8.1 The PRISM(E) Prompt Guide 46
- 8.2 Strategic roadmap: from pilot to permanent capability 48
- 8.3 Tactical roadmap: building an AI-ready practice 50
- 8.4 The Equip enablement components 50

9 Discussion and Conclusion 52
- 9.1 Answering the research question 52
- 9.2 Theoretical contribution 53
- 9.3 Practical implications 54
- 9.4 Limitations 55
- 9.5 Recommendations and future research 55
- 9.6 Personal reflection 55

Use of AI 56

References 57

Appendices 59

CHAPTER ONE

01 Introduction

The commercial telematics industry is one of the most information-dense environments in modern business. Organisations operating in this space continuously generate and receive vast quantities of market intelligence, from customer feedback and sales signals to competitive pricing data and regional market trends. For the teams responsible for translating this intelligence into actionable Go-to-Market (GtM) strategies, the challenge is not a lack of information but an excess of it. This project addresses that challenge at its source.

This report documents the progress of a graduation project conducted in partnership with Webfleet, a Bridgestone company based in Amsterdam, the Netherlands. The project investigates how GtM decision-makers at Webfleet currently capture, filter, and synthesise market insights for GtM decision-making, and proposes an AI-enabled sensemaking framework to address the information overload and workflow fragmentation that characterises current practice.

The project sits at the intersection of five bodies of literature: information overload theory, organisational sensemaking, prompt engineering, B2B knowledge creation, and situated AI. Together, these theoretical streams provide the foundation for a framework that is both practically implementable within Webfleet's existing Microsoft 365 environment and theoretically grounded in research.

CONTEXT

Webfleet - Market leader in European commercial-vehicle telematics, serving 60,000+ businesses across 180+ countries.

1.1 Webfleet, a Bridgestone company

Webfleet is Bridgestone's globally trusted fleet management solution, headquartered in Amsterdam, the Netherlands. Originally established in 1999 in Leipzig, Germany, as one of the world's earliest web-based fleet management platforms under the name Data Factory AG, the company pioneered software-as-a-service (SaaS) delivery in the fleet technology sector. Following TomTom's acquisition of Data Factory AG, the platform became known as TomTom Telematics before being acquired by Bridgestone in 2019 and rebranded as Webfleet Solutions, operating today simply as Webfleet.

Today, Webfleet is the market leader in European commercial vehicle telematics, serving over 60,000 businesses across more than 180 countries. The platform offers a comprehensive cloud-based suite covering real-time vehicle tracking, fleet optimisation, workflow management, driver-behaviour monitoring, connected navigation, and — most recently — AI-powered fleet intelligence. Webfleet operates as part of Bridgestone Mobility Solutions, and the integration of its telematics with Bridgestone's tyre-management expertise reflects Bridgestone's broader E8 Commitment.

The context most relevant to this project is Webfleet's GtM function: as a B2B technology company operating across many European and global markets, Webfleet relies heavily on its Product Marketing Managers to translate complex, fragmented market intelligence into coherent GtM strategies.

1.2 Problem context

Telematics organisations such as Webfleet operate in information-dense environments where PMMs work with fragmented inputs from multiple sources simultaneously: market research reports, customer feedback from account managers, competitive intelligence from regional sales teams, product-usage data, and strategic briefings from leadership. When information volume consistently exceeds an individual's processing capacity, decision quality deteriorates and cognitive strain increases — a condition known as information overload (Eppler et al., 2004).

At Webfleet, this manifests in several observable ways. PMMs rely heavily on individual experience and informal judgment to synthesise and prioritise market signals, resulting in inconsistent insight quality across team members and regions. The insight-synthesis process is largely manual, time-consuming, and undocumented, making it difficult to scale, audit, or improve — and there is no structured framework guiding how inputs are classified, filtered, interpreted, and communicated to downstream stakeholders such as sales teams and regional market leads.

OPPORTUNITY

AI augments human sensemaking at the point where it is most cognitively demanding — it does not replace PMM judgment.

The opportunity lies in the emerging capability of generative AI tools — specifically Microsoft Copilot, embedded within the Microsoft 365 environment that Webfleet already uses — to augment human sensemaking at the point where it is most cognitively demanding. Rather than replacing PMM judgment, AI-assisted sensemaking frameworks can structure the information-processing workflow, reduce the cognitive load associated with input classification and synthesis, and free PMMs to focus their expertise on strategic interpretation and decision-making.

1.3 Research question and sub-questions

The main research question is addressed through three sub-questions corresponding to the three phases of the project.

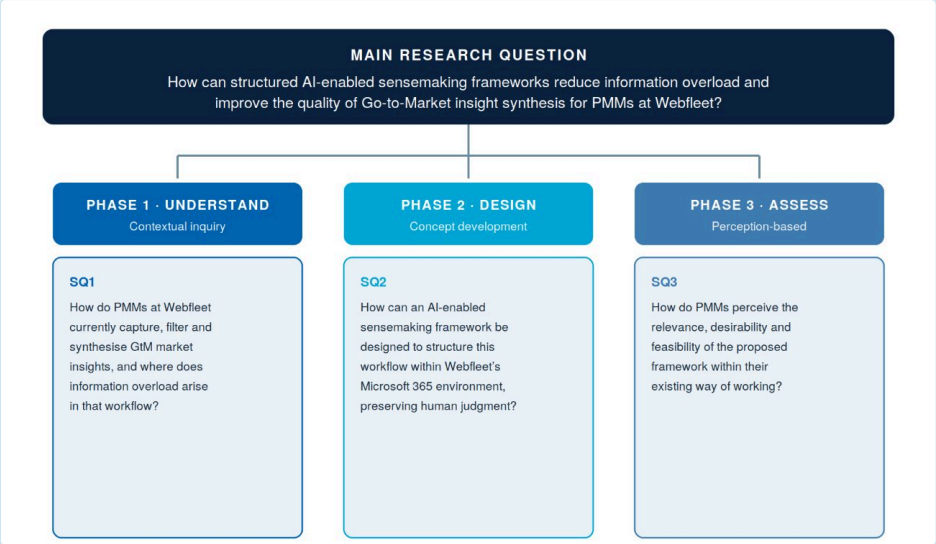


Fig. 1.1 Research question and sub-questions, mapped to the three project phases.

MARKET

Fleet management systems in use are forecast to grow from 18.1M (2024) to 30.5M by 2029 — an 11% CAGR.

1.4 Competitive landscape and market context

The European commercial vehicle telematics market is one of the most competitive and rapidly evolving technology sectors in the region. The number of fleet management systems in active use is forecast to grow at a compound annual growth rate of 11%, from 18.1 million units at the end of 2024 to 30.5 million by 2029, with market penetration in the commercial-vehicle population projected to rise from 26.7% to 42.8% over the same period. This growth intensifies the competitive pressure on Webfleet's GtM function: faster market expansion means more frequent product launches, more differentiated customer segments, and a more dynamic competitive-intelligence landscape — all of which multiply the volume of market insights that PMMs must process and synthesise.

Webfleet's primary competitors in the European market include Geotab, Samsara, Verizon Connect, Targa Telematics, and PowerFleet. According to ABI Research's 2025 commercial-telematics competitive assessment, Geotab, Samsara, PowerFleet, and Verizon Connect lead the global market across eleven innovation and implementation criteria. In the European context specifically, Targa Telematics leads by subscription volume with approximately 900,000 active subscriptions, followed by Webfleet with 752,000 at the end of 2024. Each competitor regularly releases new product features, AI capabilities, and partnership announcements — generating a continuous stream of competitive intelligence that Webfleet's PMMs must monitor, interpret, and respond to across multiple regional markets simultaneously.

CHAPTER SUMMARY → Webfleet's PMMs operate in a fast-growing, information-dense telematics market where rising launch frequency and competitive intelligence multiply the synthesis load — making information overload the core problem this project sets out to address through an AI-enabled sensemaking framework. Chapter 2 reviews the five literature streams that ground that framework and locates the gap it fills.

CHAPTER TWO

02 Literature Review

Five interconnected streams establish the theoretical foundation for the framework.

This review spans five interconnected theoretical streams: information overload theory, organisational sensemaking, prompt engineering and generative AI, B2B knowledge creation, and situated AI and digital transformation. Together, these streams establish the theoretical foundation for the framework and justify the methodological choices made in this project. All papers were selected for their direct relevance to the central problem of AI-enabled sensemaking in marketing knowledge-work contexts.

2.1 Information overload in organisational contexts

Information overload is the condition that arises when information supply consistently exceeds an individual's processing capacity, resulting in reduced decision quality, increased cognitive strain, and deteriorating performance (Eppler et al., 2004). Critically, overload is not merely a function of information volume but arises from the interaction of four factors: information characteristics (volume, ambiguity, redundancy), task complexity, organisational structures (dispersed information sharing, inadequate filtering), and individual cognitive limits.

A parallel stream reframes this condition as a form of technology-induced strain. *Techno-overload* — the way digital systems compel users to work faster and to process more information than they comfortably can — is one of the principal technostressors that erode productivity and well-being (Ragu-Nathan et al., 2008), an effect that has been linked directly to diminished role performance (Tarafdar et al., 2007). Generative AI sits on both sides of this dynamic. The same tools that promise to relieve overload can deepen it: they lower the cost of producing and surfacing information, expanding the very volume a PMM must triage, while adding a fresh layer of machine output that itself has to be checked and trusted. Read through the technostress lens, then, AI is not a straightforward remedy for overload but a force that can intensify techno-overload unless it is embedded in a structure that filters before it generates — the dual role this framework is designed to manage.

GAP

No study applies overload theory to the insight-synthesis role of B2B Product Marketing Managers.

This multi-factorial understanding is directly applicable to the PMM context at Webfleet, where all four factors are present simultaneously. Market intelligence arrives from multiple sources in varying formats and quality levels (information characteristics); GtM decision-making requires synthesising competitive, customer, and strategic dimensions at once (task complexity); insights are dispersed across regional teams, sales channels, and product groups without systematic aggregation (organisational structure); and individual PMMs perform this synthesis without structured support (cognitive limits).

Despite this foundational contribution (Eppler et al., 2004), a significant gap remains in applying information-overload theory to the specific role of Product Marketing Managers in B2B technology organisations. The reviewed literature addresses overload primarily in managerial decision-making, accounting, and general knowledge management, but does not examine the particular overload dynamics of insight-synthesis roles that sit at the intersection of multiple information streams. Furthermore, no study proposes a practical, role-specific intervention that operationalises overload reduction through structured AI-assisted workflows.

2.2 Sensemaking in organisations

Sensemaking — the process through which individuals and organisations interpret ambiguous information and construct meaning to guide action — is not a rational, linear process but an ongoing, social, and retrospective one (Weick, 1995). Sensemaking is grounded in identity, social interaction, and the need to reduce ambiguity sufficiently to enable action, even when full information is unavailable or when there is too much data. The concept of *bracketing* — the act of isolating what is relevant enough to interpret from the broader information environment — establishes that filtering is not a pre-cognitive sorting task but a meaning-making activity that requires judgment.

The design space of AI-assisted research tools for sensemaking has been investigated through a review of 11 such tools from leading HCI venues (Hao et al., 2025). That study identifies that effective AI sensemaking tools must support both divergent thinking (exploration of possibilities) and convergent thinking (narrowing toward actionable conclusions), maintain transparency of AI outputs, and preserve user agency throughout the process.

PROMPTING

Prompts function as natural-language programs that encode organisational intent and enforce reasoning patterns.

Sensemaking theory has been widely applied in organisational studies, crisis management, and strategic management (Weick, 1995), but its application to AI-assisted workflows in marketing and GtM contexts remains largely unexplored. The existing literature focuses on academic knowledge work rather than commercial marketing-intelligence workflows (Hao et al., 2025). No study applies its properties — particularly bracketing and enactment — as explicit design principles for AI tool architecture in a B2B product-marketing setting.

2.3 Prompt engineering and generative AI

The Prompt Pattern Catalog, a structured taxonomy of reusable prompt templates, establishes that prompts function as natural-language programs that encode organisational intent, structure AI outputs, and enforce reasoning patterns (White et al., 2023). Prompts with role assignment, context-setting, and output formatting consistently produce higher-quality, more reliable outputs than unstructured queries. This has been extended by reframing prompting as a form of programming — *prompt programming* — in which zero-shot prompts can outperform few-shot prompts when structured as natural-language programs with embedded narratives and meta-prompts (Reynolds et al., 2021).

The role of generative AI in augmenting knowledge work has been empirically demonstrated in a field experiment involving over 5,000 customer-support agents (Brynjolfsson et al., 2023). The findings show a 14% average productivity increase from AI tool use, with the largest gains (34%) accruing to the least experienced workers — establishing that AI tools are most transformative for those who currently lack established best practices. This is relevant to the PMM context, where junior PMMs and regional market leads without deep institutional knowledge stand to gain most from a structured AI-assisted workflow.

A further dimension comes from the investigation of EmotionPrompt, a strategy that shifts AI interactions from reactive to reflective processing (Li et al., 2025). That study demonstrates that prompts designed to encourage self-determination and reflective reasoning produce better-quality outputs in marketing decision-making, reducing cognitive biases and improving strategic insight quality.

VALUE

AI creates sustained value through enabling capabilities — the structured process, not the technology alone.

The broader applications of generative AI to business analytics have been reviewed in both marketing and innovation contexts (Agrawal et al., 2025; Cano-Marín et al., 2025), establishing that LLMs offer unprecedented capabilities for processing large datasets, identifying patterns, and generating structured knowledge summaries in marketing and innovation contexts. The prompt-engineering literature has developed robust principles for effective LLM interactions, but almost exclusively in software-development, educational, and general-productivity contexts. No study has applied prompt-pattern theory to the design of role-specific, organisation-grounded prompt libraries for marketing knowledge workers in B2B technology sectors.

2.4 AI in innovation management and B2B knowledge creation

A systematic review of AI applications in innovation management develops a taxonomy of AI capabilities across three modes: replacing human tasks, reinforcing human capabilities, and revealing new possibilities (Mariani et al., 2023). That review identifies that AI creates sustained value when it operates through *enabling capabilities* — the organisational competencies and structures that allow AI tools to be deployed effectively — rather than through technology alone. This has been extended to generative AI specifically, identifying knowledge creation, new-product development, and organisational redesign as the most significant future research directions, and highlighting the risk of strategic myopia (Peeters et al., 2024).

The relationship between AI-enabled knowledge creation and B2B marketing decision-making has been empirically examined, demonstrating that AI-enabled knowledge creation — customer knowledge, user knowledge, and external market knowledge — significantly improves marketing decision quality and firm performance (Fosso Wamba et al., 2021). The mediating mechanism is not AI technology itself but the structured knowledge-creation process it enables. There is, however, no existing study that combines AI innovation-management theory with the specific insight-synthesis challenges of B2B product-marketing roles.

SITUATED AI

Three activities — grounding, bounding, recasting — explain how AI delivers advantage when embedded in context.

2.5 Situated AI, Explainable AI (xAI), and digital transformation

The concept of Situated AI explains how organisations achieve competitive advantage by embedding AI within their specific processes, data, and strategic contexts (Enholm et al., 2023). The framework identifies three key activities: **grounding** (linking AI to real organisational problems), **bounding** (defining the limits of AI decision-making to preserve human judgment and governance), and **recasting** (adapting existing workflows to accommodate AI assistance).

The importance of explainability and transparency is established across the review (Adadi et al., 2018; and others). Collectively, the Explainable AI literature establishes that AI tools which cannot explain their reasoning to non-technical users produce trust deficits that undermine adoption — particularly in knowledge-work contexts where the consequences of acting on incorrect outputs are significant.

The organisational dimension of adoption has been addressed by two studies that both identify workforce readiness — AI literacy, confidence in outputs, and organisational culture — as the primary determinant of successful integration (Kumar et al., 2025; Ghaly et al., 2024). The first applies the Technology Acceptance Model (TAM) and Technology–Organisation–Environment (TOE) frameworks, concluding that perceived usefulness and perceived ease of use are the critical mediating variables (Kumar et al., 2025). The second establishes through qualitative case study that adoption barriers are frequently embedded in tacit workflow practices not surfaced by quantitative surveys (Ghaly et al., 2024).

A computational literature-review methodology combining machine-learning topic modelling with LLM-assisted thematic synthesis confirms the growing role of AI in supporting knowledge-synthesis tasks (Lund et al., 2024) — validating the premise that AI can meaningfully augment the synthesis and sensemaking tasks that consume significant cognitive resources in roles such as PMM.

CONVERGENCE

The streams converge on a single gap this project addresses directly.

The Situated AI framework provides a powerful account of contextually embedded AI but remains at the level of strategic theory without empirical operationalisation in specific functional contexts. The xAI literature establishes the importance of transparency, but the practical mechanisms through which explainability is embedded in organisational workflows — rather than in system architecture — remain underdeveloped. Workforce readiness is identified as critical, but no concrete design response is proposed beyond recommending training and change management (Kumar et al., 2025; Ghaly et al., 2024).

The streams converge

Information overload in knowledge-work roles is a **structural** problem, and AI only addresses it meaningfully when embedded within structured, context-specific workflows that preserve rather than bypass human judgment.

- **(Eppler et al., 2004)** & **(Ragu-Nathan et al., 2008)** diagnose the problem — overload as a structural and techno-stress condition.
- **(Weick, 1995)** establishes that sensemaking is a human interpretive act AI must augment, not automate.
- **(White et al., 2023)** & **(Reynolds et al., 2021)** provide the operational mechanism — structured prompt engineering.
- **(Fosso Wamba et al., 2021)** & **(Mariani et al., 2023)** confirm value lies in the structured process, not the technology.
- (Enholm et al., 2023) supply the strategic rationale: value comes from being situated in context.
- **(Kumar et al., 2025)** & **(Ghaly et al., 2024)** add the final condition — trust and workforce readiness are prerequisites, not afterthoughts.

Looked together, these reveal a clear gap: no existing study has designed or assessed an AI-enabled sensemaking framework for the GtM insight-synthesis workflow of B2B Product Marketing Managers in a telematics context. This project addresses that gap directly.

CHAPTER SUMMARY → Five streams converge on one conclusion: overload is structural, sensemaking is a human act AI must augment, structured prompting is the mechanism, value lies in the process, and situated fit plus workforce readiness are the conditions — yet no study joins them for the PMM role. Chapter 3 introduces PRISM(E), the framework built to close that gap.

CHAPTER THREE

03 The PRISM(E) Framework

As a prism refracts light into a spectrum, PRISM(E) refracts fragmented intelligence into structured insight.

PRISM(E) is an AI-enabled sensemaking framework operationalised through structured prompts and decision logic, designed to support PMMs at Webfleet in processing fragmented market inputs into GtM-ready insights. It stands for **Prioritise, Reduce, Interpret, Synthesise, Mobilise,** and **Equip**. The first five steps constitute the sequential insight workflow; the sixth — Equip — is an organisational enablement layer that runs alongside all five and addresses the AI literacy and readiness conditions necessary for adoption.

The name reflects its function: as a prism refracts undifferentiated light into a structured, usable spectrum, PRISM(E) takes fragmented market intelligence and refracts it into a structured, prioritised, explainable GtM insight. The framework does not eliminate or compress information — it transforms it into a form usable for strategic decision-making.

DEVELOPMENT

PRISM(E) was not adopted from an existing model — it was constructed by translating the literature and the Phase-1 evidence into a workflow.

3.1 Developing the framework

PRISM(E) is an original framework developed for this project rather than an existing model applied to a new setting. No single source in the literature offered a workflow for the problem at hand, so it was built by synthesising the five theoretical streams of Chapter 2 with the design targets from the contextual inquiry (Chapter 5). The development followed a deliberate, traceable logic in which every stage answers a requirement drawn from the literature, so the framework is evidence-led rather than intuition-led.

The starting point was the diagnosis. Information-overload theory and the technostress literature establish that overload is structural and multi-factorial (Eppler et al., 2004; Ragu-Nathan et al., 2008), which set the first requirement: any intervention must *filter before it interprets*. Filtering and interpretation are themselves human meaning-making acts (Weick, 1995), setting the second: the workflow must *preserve a human gate* at every consequential step rather than automate it away. The prompt-engineering stream (White et al., 2023; Reynolds et al., 2021) supplied the operational mechanism, the knowledge-creation stream (Fosso Wamba et al., 2021) located value in the structured process, and the situated-AI and adoption streams (Enholm et al., 2023; Kumar et al., 2025) added the conditions of contextual fit and workforce readiness. Each requirement was translated into a design principle and then sequenced into a workflow that mirrors how PMMs already reason — from raw inputs to a shared, distributed insight.

1. **Diagnose.** Synthesise the five streams and the Phase-1 evidence into a single problem statement.
2. **Derive principles.** Translate each stream into a concrete design requirement — filter, bracket, interpret, verify, distribute, enable.
3. **Sequence.** Order the requirements into a workflow that follows the PMM's own path from raw signal to shared insight.
4. **Name.** Encode the workflow as the PRISM(E) mnemonic so it is memorable and adoptable — itself an Equip consideration.
5. **Ground.** Situate every stage inside Webfleet's existing Microsoft 365 environment and governance.
6. **Iterate.** Refine this conceptual v1 against the interview findings and co-creation into the validated v2 (Chapters 6–7).

The name carries the core idea. As a prism refracts undifferentiated light into a structured spectrum, the framework refracts fragmented market intelligence into a structured, prioritised, explainable insight; the verbs map one-to-one onto the workflow stages. A sixth element, Equip, was added once the adoption literature made clear that a sound structure still fails without the readiness conditions to sustain it. The five stages and the Equip layer are described next.

TOOLING

Tool choice follows the Situated-AI principle: value arises from contextual fit, not raw capability.

Why inside Microsoft 365 Copilot

All five workflow steps run inside Microsoft Copilot within Webfleet's own M365 tenant — a choice grounded in four contextually specific reasons:

- **Data residency.** Data stays within Webfleet's environment, addressing GDPR compliance.
- **No new procurement.** No infrastructure purchase is required, removing the readiness barrier.
- **Native integration.** Copilot connects to SharePoint, Teams, and Power Automate — enabling the automation in P and M.
- **Legal defensibility.** Microsoft's enterprise Data Processing Agreement covers all interactions under the existing licence.

3.2 P — Prioritise

Prioritise addresses the most acute dimension of overload: the cognitive burden of processing all incoming information at equal priority before any filtering. All inputs for a week are linked in a SharePoint intake folder, auto-sorted by Power Automate into category subfolders. Copilot then classifies each input and ranks it by relevance to the current GtM priority, scoring each 1–3 on urgency and source reliability. The PMM begins the week with a structured, ranked shortlist rather than an undifferentiated pile of documents — operationalising the *grounding* activity of Situated AI (Enholm et al., 2023).

3.3 R — Reduce

Reduce addresses the concept of *bracketing* (Weick, 1995). From the ranked shortlist, Copilot identifies the two to three inputs most directly relevant to the current decision and provides one sentence of reasoning for each. The PMM then confirms or adjusts. This confirmation is the most important design decision in the stage: it is not optional and cannot be automated, because relevance is context-dependent in ways AI cannot fully assess without the PMM's tacit knowledge. Human-agency preservation is a core design principle that this requirement directly operationalises (Hao et al., 2025).

HUMAN GATE

Interpret produces a draft, not a conclusion. Synthesise is where human judgment does its work.

3.4 I — Interpret

Interpret is the AI's primary contribution. Copilot receives the confirmed shortlist and generates a structured insight draft using a B2B-analyst synthesis prompt. The prompt assigns Copilot the role of a B2B marketing analyst in a telematics company, grounds it in Webfleet's GtM context, and specifies a three-part output: what is happening, why it matters for GtM strategy, and what action or watch-point is recommended. This structure — role assignment, context-setting, output formatting — consistently outperforms unstructured queries (White et al., 2023). The Interpret step produces a draft, not a conclusion: the PMM does not act on Copilot's output here.

3.5 S — Synthesise

Synthesise is the PMM's primary contribution and the most intellectually demanding stage. The PMM reads the AI draft critically, adds Webfleet-specific context Copilot cannot access — a recent conversation with a key account manager, a known roadmap constraint, a regional competitor dynamic — and rewrites the strategic so-what in their own words. Copilot then formats the approved insight. Before finalising, the PMM completes a five-question audit gate embedded in the SharePoint template:

THE SYNTHESISE AUDIT GATE

1. Was the correct tool used for this data-sensitivity level?
2. Does the insight contain personal data that must be removed before sharing?
3. Can I name the source of every claim?
4. Has the AI output been reviewed and edited?
5. Would I be comfortable if our legal team reviewed how this was produced?

All five must be answered affirmatively before the insight proceeds. This operationalises the xAI literature's transparency requirements not as a system property but as a human workflow checkpoint — consistent with the *bounding* activity of Situated AI (Enholm et al., 2023) and the *enactment* principle of sensemaking (Weick, 1995).

COMPOUNDING

The SharePoint archive turns one-off insights into a searchable repository the team builds on.

3.6 M — Mobilise

Mobilise distributes the approved insight to the right stakeholders in the right format. Copilot drafts three channel-specific versions: a Teams summary for sales, an executive brief for leadership, and a regional note for market leads. Power Automate pushes these on the PMM's approval, and the full insight is archived in the SharePoint insight library.

The knowledge-base dimension is its most strategically significant feature. Over time, the archive becomes a searchable repository that future Copilot queries can draw from, letting the team build on previous insights rather than starting from scratch each cycle. This compounding knowledge effect produces the largest long-term productivity gains (Brynjolfsson et al., 2023), and value is sustained when outputs are systematically stored rather than discarded (Mariani et al., 2023).

3.7 E — Equip

Equip is the sixth and structurally distinct component. Unlike the five sequential steps, it is an organisational enablement layer that runs alongside the entire framework, addressing the human and organisational conditions necessary for PRISM(E) to be adopted sustainably. It covers four dimensions identified across the adoption literature: **AI literacy** (knowing how and when to use AI), **confidence** (feeling safe to use AI-assisted outputs in accountable decisions), **trust calibration** (knowing when to override AI), and **prompt fluency** (adapting the templates when defaults do not fit).

Workforce readiness is the primary determinant of successful adoption — ahead of technological capability (Ghaly et al., 2024). Confidence is the critical mediating variable (Kumar et al., 2025), and structured guidance at the point of introduction separates teams that adopt AI from those that do not (Brynjolfsson et al., 2023).

FIGURE 3.1

Five sequential stages, with Equip running beneath them as the enablement layer.

The conceptual framework, at a glance.

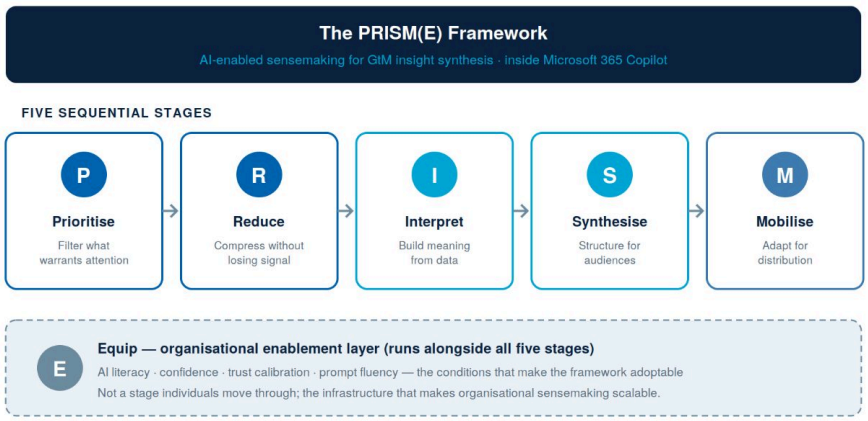


Fig. 3.1 PRISM(E) framework summary (conceptual version).

CHAPTER SUMMARY → PRISM(E) refracts fragmented market intelligence into structured, prioritised, explainable insight through five sequential stages — Prioritise, Reduce, Interpret, Synthesise, Mobilise — with Equip running beneath as an enablement layer. As presented here it is a literature-derived conceptual version. Chapter 4 sets out the qualitative, design-led methodology used to test and refine it.

CHAPTER FOUR

04 Methodology

This project adopts a qualitative, design-led research approach grounded in design science research (DSR). The research design is sequential and iterative, structured across three phases that correspond directly to the three sub-questions. Each methodological choice is justified by evidence from the literature review rather than imported from an external tradition.

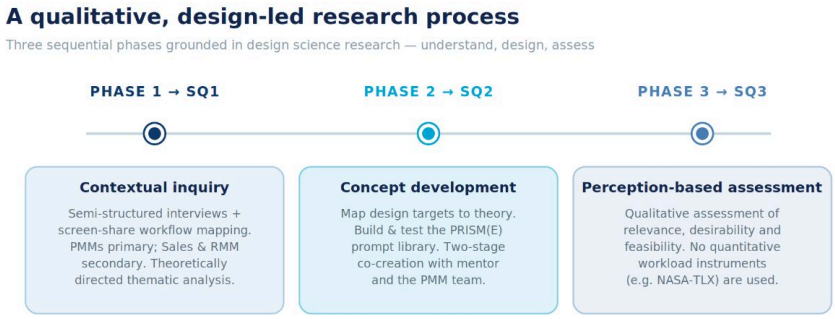


Fig. 4.1 Three-phase methodology mapped to the sub-questions — understand, design, assess.

STANCE

Constructivist: knowledge is contextually constructed through researcher, participants, and the Webfleet environment.

4.1 Research approach

The project is situated within a constructivist epistemological tradition, consistent with the sensemaking paradigm (Weick, 1995). Knowledge is understood as contextually constructed through the interaction between researcher, participants, and the organisational environment at Webfleet — not as objective truth to be extracted through measurement. This stance is appropriate for a problem that is fundamentally about how individuals interpret and make meaning from complex information environments.

The design-led approach is justified by the nature of the output: a practical artefact — a framework and prompt library, with a roadmap to implement it and a workshop to improve AI literacy. AI tools in innovation-management contexts must be designed through empirical grounding in real organisational problems rather than purely theoretical development (Mariani et al., 2023). The three-phase structure — understand, design, assess — reflects this directly.

4.2 Contextual inquiry (Phase 1)

This phase addresses SQ1 through semi-structured interviews and workflow mapping with PMMs and secondary stakeholders. Semi-structured interviews are appropriate because sensemaking accounts are retrospective narratives that emerge through conversation (Weick, 1995). The interview guide contains six questions per stakeholder group across three phases — opening, core, and closing — with parallel-structure questions so findings from different roles can be directly compared and triangulated.

Adoption barriers are frequently embedded in tacit workflow practices not surfaced by verbal accounts alone (Ghaly et al., 2024). Participants are asked to share their screen and narrate their current synthesis process in real time. Data is analysed using theoretically directed thematic analysis — inductive on a first pass, deductive on a second against the four overload dimensions (Eppler et al., 2004) and sensemaking properties (Weick, 1995). The output is three to five design targets that directly inform Phase 2.

DESIGN WITH

Co-design preserves human agency and builds the trust conditions adoption depends on.

4.3 Concept development (Phase 2)

This phase addresses SQ2 through iterative design, prompt engineering, and co-design sessions. The framework is developed by mapping the Phase 1 design targets onto the theoretical constructs from the literature review, so every design decision traces back to an empirical or theoretical source rather than researcher intuition.

The prompt library is designed using established prompt-pattern principles — role assignment, context-setting, output formatting (White et al., 2023) — and tested against realistic Webfleet scenarios before being presented to participants. Prompt effectiveness can be evaluated through systematic output comparison, which informs the testing protocol (Reynolds et al., 2021).

The use of co-design is supported by the human-AI augmentation principle (Hao et al., 2025): AI sensemaking tools must be designed *with*, not merely *for*, the knowledge workers who will use them — preserving agency and building the trust conditions identified as essential (Kumar et al., 2025).

4.4 Evaluation (Phase 3)

Phase 3 addresses SQ3, framed deliberately as a perception-based question rather than an evaluative one. It does not measure effectiveness through controlled experiments or quantitative workload instruments, which would be inappropriate for a 100-working-day project and inconsistent with the constructivist stance. Instead it assesses how PMMs perceive the relevance, desirability, and feasibility of the framework, mapping to the perceived-usefulness, perceived-ease-of-use, and organisational-readiness constructs (Kumar et al., 2025).

Because the framework cannot be fully lived-in within the thesis window, assessment is conducted through co-creation rather than prototype trialling. Practitioners engage with the framework and prompt library, reflect on their fit, and their input feeds iterative refinement between sessions, justified by the finding that quantitative metrics alone cannot surface the trust and governance concerns that determine adoption (Ghaly et al., 2024).

Validation is split by stakeholder weight. **PMMs, the primary group, are validated empirically** — they reviewed the finished guide in a workshop and responded to it directly. For the secondary groups (RMMs and Sales), no live review was feasible within the project window, so they are assessed for *anticipated fit*: two synthetic personas — one per group, built as composites of the role scope and pain points participants described in their own interviews — are walked against the guide's relevant sections. The personas are treated strictly as analytical instruments, not endorsement; every participant quotation used in validation is verbatim, and one Sales participant additionally supplied a written response used as direct evidence of desirability.

GOVERNANCE

All sessions used Copilot inside Webfleet's tenant — the bounding activity, demonstrated in the research design itself.

4.5 Sampling strategy

Purposive sampling is used throughout all three phases. Information overload is a context-specific phenomenon (Eppler et al., 2004), contingent on particular combinations of information characteristics, task complexity, and organisational role. The sample must therefore consist of participants who work within Webfleet's GtM ecosystem and have direct experience of fragmented market-intelligence synthesis. The multi-stakeholder design — PMMs as primary, sales and regional leads as secondary — reflects the finding that studying AI value creation across multiple roles is essential for understanding how AI embeds at different levels of an organisation (Enholm et al., 2023). The same interviews serve a second purpose in Phase 3: for the secondary groups they supply the role scope and pain points from which the validation personas are composited, so the anticipated-fit assessment remains grounded in primary data rather than assumption.

4.6 Ethical considerations and data governance

The data-governance implications of AI tool use within the research are addressed through the exclusive use of Microsoft Copilot within Webfleet's existing M365 tenant for all feedback sessions. This ensures that no commercially sensitive market intelligence used as task input leaves the Webfleet environment. The constraint is not only procedural: it operationalises the *bounding* activity of Situated AI (Enholm et al., 2023) directly within the research design itself, ensuring the governance layer built into PRISM(E) is authentically demonstrated during the assessment sessions.

CHAPTER SUMMARY → A qualitative, design-led DSR method runs across three phases — understand, design, assess — using contextual inquiry, co-creation, and perception-based validation split by stakeholder weight, with purposive sampling and Copilot kept inside Webfleet's M365 tenant for governance. Chapter 5 reports the findings of the Phase 1 contextual inquiry.

CHAPTER FIVE

05 Findings from the Contextual Inquiry

Ten interviews: five PMMs, two RMMs, three from Sales.

The empirical phase produced findings from ten interviews: five Product Marketing Managers (PMM1–PMM5), two Regional Market Managers (RMM1–RMM2), and three members of the Sales team (Sales1–Sales3).

PMMs were treated as the primary group because GtM insight synthesis sits at the centre of their role; RMMs and Sales were secondary — the regions translating central output into local markets, and Sales carrying it to the customer. Rather than report each group in isolation, this chapter reads the three groups together. A single framework has to serve all of them at once, so the useful question is not what distinguishes a PMM from a sales rep, but where their needs converge, where they diverge, and what that pattern asks of the design.

ANALYSIS

Two coding cycles — inductive, then deductive against the overload and sensemaking constructs.

"We're not the experts at writing good prompts — that's exactly what we'd want centralised."

RMM2

5.1 Analysing the interview data

All ten interviews were conducted in English, recorded with participant consent, and transcribed verbatim. The transcripts were analysed through a theoretically directed thematic analysis following a structured six-phase thematic-analysis procedure. Coding proceeded in two cycles. The first was an inductive, open-coding pass that stayed close to participants' own language, labelling each meaningful segment with a descriptive code. The second was a deductive pass that mapped those codes against the constructs from the literature review — the four overload factors (Eppler et al., 2004) and the sensemaking properties of organisational meaning-making (Weick, 1995) — so the empirical material could be read directly against the theory the framework is built on.

Codes were then clustered into candidate themes and tested across the three role groups. A pattern was retained as a finding only where it was named explicitly by several PMMs or corroborated by at least one secondary group, which guarded against over-reading a single account; the screen-share walkthroughs were used to check stated practice against observed behaviour. Because the interview guide was parallel-structured across roles, responses could be compared directly and triangulated. The analysis converged on the shared standards, frictions, and four cross-cutting gaps reported in the rest of this chapter, with representative quotations used throughout to anchor each theme in participants' own words.

SAMPLE

Ten participants across three role groups, spanning the GtM chain from central synthesis to local activation.

5.2 Respondent characteristics

The ten participants span the GtM chain from central insight synthesis to local activation. PMMs form the analytical core; Regional Market Managers and Sales are secondary, representing how central output is localised and carried to the customer.

ID	ROLE / POSITION	GROUP	PRIMARY FOCUS AND VANTAGE POINT
PRIMARY · PRODUCT MARKETING — INSIGHT SYNTHESIS			
PMM1	Product Marketing Manager	Primary	Cross-domain portfolio; the weekly problem is deciding what matters across many fragmented topics.
PMM2	Product Marketing Manager	Primary	One technically dense domain with no internal knowledge holders; uses AI mainly to translate and onboard.
PMM3	Product Marketing Manager	Primary	Expansion-minded; treats AI as a way to raise output volume; trusts win/loss CRM data most.
PMM4	Product Marketing Manager	Primary	Hypothesis-driven and prompt-conscious; the only interviewee building personal AI agents.
PMM5	Product Marketing Manager	Primary	Reasons at team-system level (~30 concurrent products); tiers trust in AI by what is at stake.
SECONDARY · REGIONAL MARKETS			
RMM1	Regional Market Manager	Secondary	Localises central content; competitor monitoring is the main load trigger.
RMM2	Regional Market Manager	Secondary	Strategy-driven prioritisation; load peaks during Sept–Oct annual planning.
SECONDARY · SALES			
Sales1	Sales	Secondary	Activation edge; loses ~20 min per lead to manual research; Copilot is the primary tool.
Sales2	Sales	Secondary	Activation; wants fill-in-the-blank prompt templates to contextualise market data per prospect.
Sales3	Sales	Secondary	Activation; focused on prioritisation and conversion; proposes specialised agents across the workflow.

Table 5.1 Respondent characteristics — the ten interview participants across three role groups, spanning the GtM chain from central insight synthesis to local activation.

PATTERN

Overload is not one phenomenon — breadth, depth, and structure each produce a different version of it.

"Too many domains compete for attention at once; the weekly task is deciding what even matters."

PMM1

5.3 The stakeholder sample

The five PMMs form the analytical core and differ mainly in vantage point. PMM1 works across many domains at once, so the weekly problem is deciding what matters across fragmented topics. PMM2 owns a single technically dense domain with no internal knowledge holders, making AI valuable above all as a translation and onboarding tool. PMM3 is the most expansion-minded, treating AI as a way to raise output volume. PMM4 is the most methodologically deliberate — hypothesis-driven, conscious of prompt design, the only interviewee actively building personal AI agents. PMM5 reasons at the level of the team system, preoccupied with how a structure absorbs ~30 concurrent products without collapsing.

The two RMMs sit between central marketing and local markets. Both confirmed central content is generally well produced and located the problem in the final localisation step. The three Sales interviewees represent the activation edge, where an insight either converts into a customer conversation or does not; all three rely on Copilot as their primary tool.

5.4 How overload is experienced

Overload did not present as a single phenomenon. For PMM1 it is breadth — too many domains competing for attention. For PMM2 it is depth — getting on top of one complex field with no internal expertise. PMM5 experiences it structurally, as a question of designing a team-level system. RMM1's biggest trigger is competitor monitoring; RMM2's load spikes during annual planning in September–October. Sales feel it as fragmentation rather than volume; as Sales3 put it, the problem is prioritisation and conversion into something usable.

What unifies these accounts is the absence of a shared filter. No team-level criteria exist for deciding what counts as priority versus noise. The one place filtering works cleanly is where it is anchored to a declared strategic focus: RMM2's prioritisation is strategy-driven — if construction is the focus vertical for the year, passenger-transport data is deprioritised regardless of how much is available. That contrast points directly at the design target.

“An exhaustive Excel table of accurate figures is not an insight.”

PMM5

5.5 A shared standard for a good insight

Across all ten interviews a shared standard emerged: a good insight is evidence-based — anchored to a clear number or data point, structured, and quick to grasp; a weak one is too long and demands too much effort to extract the takeaway. The rule that an insight must carry a numerical anchor before it can be shared externally held without exception across the five PMMs. Within that bar, emphases shift by role: PMM1 and PMM3 stress actionability, PMM2 the numerical anchor, PMM3 the why behind a signal, PMM4 frames insight as the validation of a hypothesis, and PMM5 defines it as communication that explains itself.

Beneath all of this lies a single distinction the team does not always draw before sharing: **data is not insight**. As PMM5 put it, an exhaustive Excel table of accurate figures is not an insight; an insight is data given structure — a trend line, a takeaway, and a message. The shared definition is less about accuracy than about the work of turning accurate material into something that travels.

5.6 Trust in AI

Trust in AI was universal in form and conditional in every case. All interviewees agreed that AI output requires human review before any stakeholder use, and no one trusts the tools to work unsupervised on anything consequential — framed as the correct instinct rather than a limitation. What varies is calibration to the stakes \u2014 PMM1 is the most cautious, having once found AI-generated compliance content built on outdated regulations:

> *"The moment you catch the AI citing rules that no longer apply, you learn never to let it run unsupervised on anything that matters".*
>
> PMM1 PRODUCT MARKETING MANAGER

PMM4 wants the reasoning chain made visible; PMM5 wants a tool's known weaknesses declared in advance and tiers her trust by how much is riding on the result.

The secondary groups confirm the same reflex. Both RMMs verify sources before publishing; Sales treat AI-generated statistics as directional ranges rather than definitive claims. Trust, across the board, is procedural rather than blanket.

HIDDEN COST

Validation takes 3–4× the effort of extraction; treatment crowds out the thinking.

"We spend most of our analytical time formatting and cleaning data, not actually thinking."

PMM4

5.7 Data treatment

Two findings about the mechanics of working with data recurred across roles. The first is that **treatment is crowding out analysis**. PMM4 named it directly: PMMs spend the majority of their analytical time on formatting and cleaning rather than thinking. PMM3 framed it as structural rather than a skill gap, with validation taking three to four times the effort of extraction — the root being that data tools report inconsistent numbers across sources, and reconciling them consumes time out of all proportion to the value.

The second is that **sources are not equally reliable**, and the difference is systematic enough to act on. PMM5 described consistent, team-wide over-promising in sales forecasts — a structural source-reliability problem, not an individual one. PMM3, by contrast, treats win/loss CRM data as the most commercially grounded signal available. Together these justify weighting source reliability explicitly inside the Interpret stage rather than treating every input as equivalent.

5.8 Structural friction beneath the information problem

Several findings point past information handling to organisational structure. PMMs are repeatedly brought into projects late, creating downstream time pressure no tool can fully resolve. For RMMs the friction is the localisation gap: cultural nuance, channel simplification for resellers, and local regulatory detail. On the Sales side the unsolved problem is discoverability — well-organised PMM material stays underused because not all agents navigate the content repository effectively, and Sales1 loses roughly twenty minutes per lead to manual research. Where the organisation leaves a gap, AI sometimes fills it: PMM2 used it as an onboarding tool, translating dense documentation into accessible language. The framework cannot redesign these processes, but it has to stay useful inside them.

THE ASK

Not more AI access — a shared, reusable prompt scaffold.

"I don't need more access to AI. I need a ready-made prompt I can drop a prospect into."

SALES3

5.9 A shared demand for prompting support

Prompting quality maps directly onto output quality, and the appetite for help with it surfaced independently in every group. PMM4 is already building personal AI agents. RMM2 asked outright for centralised prompting support, noting the team are not the experts at writing good prompts. Sales2 and Sales3 both requested pre-built, fill-in-the-blank prompt templates that contextualise generic market data to a specific prospect, and Sales3 went further, setting out an ideal product-brief format and proposing specialised agents across the workflow. The signal is consistent: the demand is not for more AI access but for a shared, reusable prompt scaffold — precisely what the prompt library answers.

5.10 Four gaps and the design targets

Read together, the stakeholder findings converge on four cross-cutting gaps the framework had to close. Each maps onto a stage of PRISM(E) and gives the design an evidence-based target rather than a generic ambition.

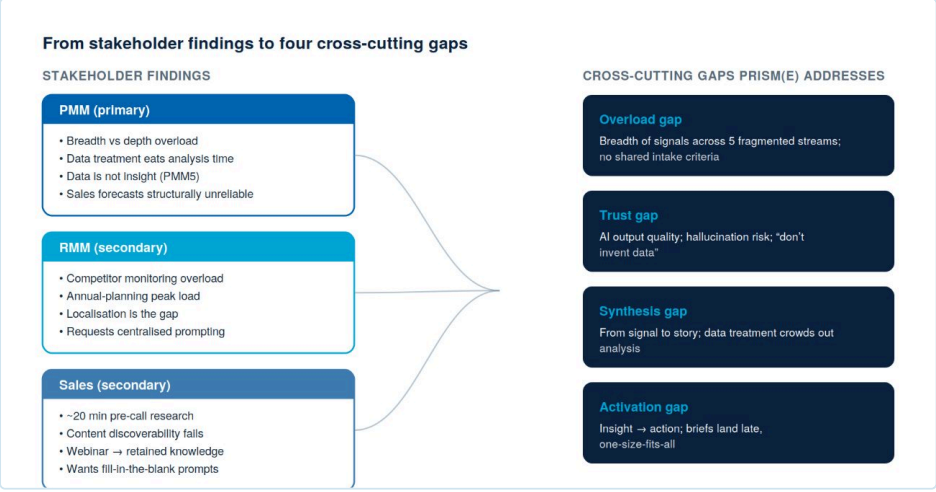


Fig. 5.1 Thematic findings across stakeholders, consolidated into the four gaps the framework addresses.

DESIGN TARGETS

Each gap becomes a specific, evidence-based target carried into Phase 2.

From four gaps to four design targets.

GAP 01 · PRIORITISE + REDUCE

Overload

Five fragmented signal streams with no shared intake criteria to govern them.

Target: surface the few signals that warrant attention against a stated focus.

GAP 02 · SYNTHESISE

Trust

Hallucination risk and the firm, shared rule never to invent data.

Target: a verification discipline built into the workflow itself.

GAP 03 · INTERPRET + SYNTHESISE

Synthesis

Data treatment crowds out analysis; data is shared as though it were already insight.

Target: move reliably from a raw signal to a structured, self-explanatory story.

GAP 04 · MOBILISE

Activation

Briefs that arrive late and in a one-size-fits-all form.

Target: the right insight, for the right audience, in the right format.

CHAPTER SUMMARY →

Across ten interviews, overload at Webfleet is consistently a fragmentation and validation burden rather than a volume problem; the inquiry surfaced shared insight standards, conditional trust in AI, and four gaps — overload, trust, synthesis, and activation — that became the Phase 2 design targets. Chapter 6 shows how these refined the conceptual framework into its second version and shaped the prompt library.

CHAPTER SIX

06 Refined Framework and Design

The six elements held; what the interviews changed was the content of each stage.

In Phase 2, the design targets from the contextual inquiry were turned into two connected artefacts: a refined version of PRISM(E), and a prompt guide that makes each stage usable inside Microsoft Copilot.

6.1 From design targets to a refined framework

The conceptual framework in Chapter 3 was derived from literature. Phase 1 tested it against how the work is actually done, and the findings did not overturn the structure so much as sharpen it. The six elements held, and no stage was removed or reordered — itself a finding worth stating plainly: the sequence reflects how practitioners already reason, even where they lack a shared vocabulary for it. What the interviews changed was the content of each stage, and the Equip layer expanded from a single enablement idea into a three-part structure.

GAP · PHASE 1	DESIGN TARGET	STAGE
Overload BREADTH OF SIGNALS	Surface the few signals that warrant attention against a stated focus.	P R
Trust AI OUTPUT QUALITY	Build a verification discipline into the workflow itself.	S
Synthesis SIGNAL TO STORY	Move reliably from a raw signal to a structured, self-explanatory story.	I S
Activation INSIGHT TO ACTION	The right insight, for the right audience, in the right format.	M

Fig. 6.1 The four gaps identified in Phase 1 became the design targets carried into Phase 2, each addressed by specific PRISM(E) stages.

V1 → V2

Every refinement traces back to a specific finding from the inquiry.

6.2 Stage-by-stage refinements

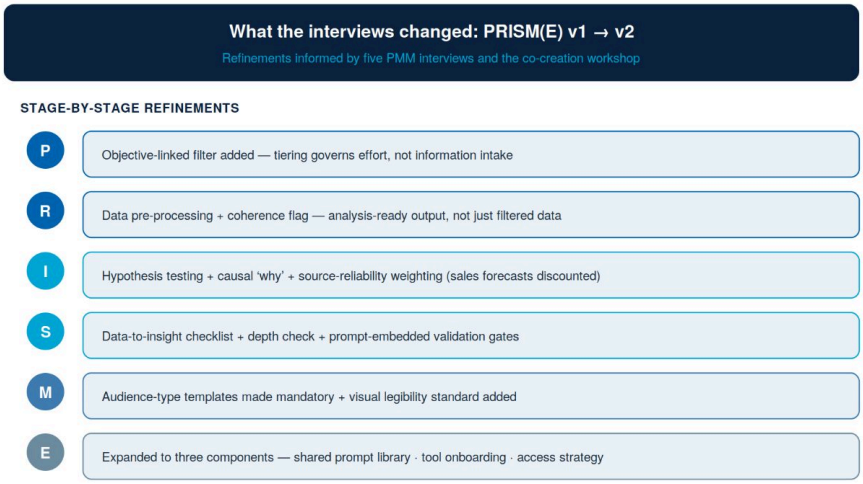


Fig. 6.2 Stage-by-stage refinements from the conceptual to the refined framework.

P **Prioritise** gained an objective-linked filter — relevance is tied to the current GtM question, because tiering governs effort but not information intake.

R **Reduce** gained a data pre-processing step and a coherence flag, surfacing conflicting numbers before interpretation begins.

I **Interpret** gained hypothesis testing, a required causal "why", and source-reliability weighting — where the sales-forecast finding lives.

S **Synthesise** gained a data-to-insight checklist, a depth-appropriateness check, and validation checkpoints embeddable inside the prompt.

M **Mobilise** made audience-type templates mandatory and added a visual-legibility standard so an insight communicates without verbal support.

E **Equip** expanded into three components: a shared prompt library, tool-specific onboarding, and an organisational access strategy.

THREE PROMPT TYPES

A copy-paste prompt, a template to fill in, and a short statement of prompt logic — for every stage.

6.3 Designing the prompt library

The prompt library is the practical mechanism that makes the framework usable day to day. It was designed using established prompt-pattern principles — role assignment, context-setting, output formatting (White et al., 2023) — and tested against realistic Webfleet scenarios before being presented to participants. Every stage offers three prompt types so practitioners of different confidence levels can all use it: a **copy-paste prompt** for immediate use, a **template prompt** to fill in for a specific context, and a short **statement of prompt logic** explaining why the prompt is built the way it is. Each prompt traces back to a finding, so the library reads as a direct response to the pain points rather than a generic set of templates.

The library is organised both by PRISM(E) stage and by role, with dedicated sections for PMMs, RMMs, and Sales, because the inquiry showed the same stage serves different tasks in different roles. A session-level context prompt sets the role, tone, insight-quality standard, and trust and verification rules once, so the individual stage prompts inherit a consistent frame. This directly answers the requests from Sales and RMM2 for fill-in-the-blank prompts that reduce the prompting skill required while keeping output quality high.

> GOVERNING PRINCIPLE
>
> AI is a **synthesis and reduction tool**, not an autonomous decision-maker — a principle stated throughout the guide.

WITH, NOT FOR

Workshop output treated as data, not endorsement.

6.4 Co-creation with practitioners

The framework and prompt library were refined *with* practitioners rather than only *for* them, consistent with the human-AI augmentation principle (Hao et al., 2025). Co-creation was run in two stages.

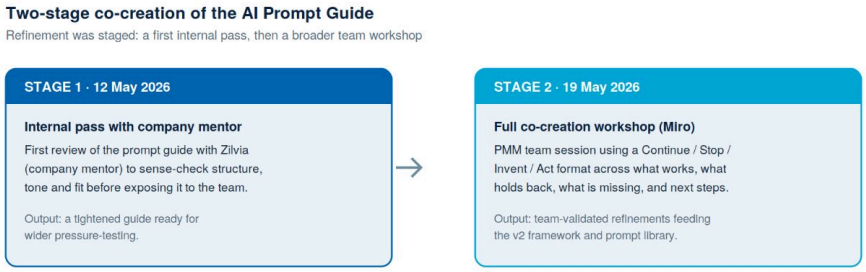


Fig. 6.3 Two-stage co-creation process behind the refined framework.

The first stage, on 12 May 2026, was an internal pass with the company mentor — sense-checking the structure, tone, and organisational fit before exposing the guide to the wider team. The second, on 19 May 2026, was a full co-creation workshop with the PMM team on Miro using a Continue / Stop / Invent / Act format. Treating the output as data rather than endorsement, its contributions fed the refinements above — particularly the move to mandatory audience templates, the depth-appropriateness check, and the expansion of Equip. In keeping with the perception-based stance, the sessions informed the design; they did not constitute a formal evaluation of effectiveness, which is reserved for the validation phase.

CHAPTER SUMMARY → Phase 2 kept all six elements but sharpened the content of each stage, expanded Equip into three components, and produced a role- and stage-organised prompt library in three forms — co-created with practitioners rather than designed only for them. Chapter 7 validates the resulting prompt guide against relevance, desirability, and feasibility.

CHAPTER SEVEN

07 Validation of the Prompt Guide

Perception-based validation: relevance, desirability, feasibility — empirical for PMMs, anticipated-fit for the secondary groups.

This chapter validates the central deliverable — the PRISM(E) Prompt Guide — against the third sub-question (SQ3): how the artefact is perceived in terms of relevance, desirability, and feasibility.

Consistent with the qualitative orientation of the study, validation is perception-based rather than measurement-based. It does not quantify time saved or output quality; it asks whether the guide addresses problems the stakeholders recognise as their own, whether they would want to use it, and whether it fits the constraints of their actual workflow.

ASYMMETRY

PMMs are the primary stakeholders and validate empirically; RMMs and Sales are secondary and are tested for anticipated fit.

7.1 Validation approach and rationale

The guide is the operational layer of a three-part structure: PRISM(E) is the sensemaking framework, the prompt guide operationalises it inside Microsoft Copilot, and the implementation roadmap situates adoption over time. The framework and roadmap are evaluated elsewhere; the purpose here is narrower — to establish whether the prompt guide is relevant, desirable, and feasible for the people whose work it is designed to support.

7.1.1 Primary and secondary stakeholders

The guide serves three internal functions, but they do not carry equal weight in the design rationale. **Product Marketing Managers (PMMs) are the primary stakeholders**: the framework was shaped around their workflow, and the bulk of the guide is PMM-specific. **Regional Market Managers (RMMs) and Sales are secondary**: they consume and act on PMM output, and the guide extends to them because the same sensemaking logic applies to their downstream work. This asymmetry shapes how each group is validated. PMM validation is empirical — PMMs reviewed the finished guide directly. For RMMs and Sales, no live review of the finished guide was possible within the project window, so their validation is an *anticipated-fit* assessment: the guide's sections are tested against the documented needs captured in the research-stage interviews, with one Sales participant additionally providing direct written evidence.

7.1.2 Persona-based method and its boundaries

To assess anticipated fit for the secondary stakeholders, two synthetic personas were constructed — one for RMMs, one for Sales. Each is a composite built from two sources only: the role scope participants described in their own interviews, and the pain points they reported. Two boundaries are stated plainly. First, the personas are *analytical constructs, not stand-ins for endorsement* — they walk the guide's prompts against realistic needs but do not claim the individuals behind them approved the artefact. Second, every quotation in this chapter is *verbatim* from the transcripts; where a finding cannot be supported by a real statement, it is presented as analysis rather than evidence.

PRIMARY

The strongest signal in the project: the primary group reacting to the finished artefact, not an early concept.

"The problem is validating the data — if it's true or not."

PMM · TECHNICAL PORTFOLIO

7.2 Primary validation: Product Marketing Managers

The PMM interviews confirm the study's central finding: overload at Webfleet is a **fragmentation and validation burden, not a volume problem**. The expensive work is not finding data but trusting it. One PMM put the cost squarely on verification rather than extraction, describing spending several times longer justifying a report than producing it:

> *"The problem is validating the data — if it's true or not."*
>
> PMM · TECHNICAL PRODUCT PORTFOLIO

A second located the difficulty in the absence of an organising structure, and a third pointed directly at the prompting gap the guide is built to close:

> *"I wouldn't say it's overwhelming as long as you come in with a very clear structure how you want to build it."*
>
> PMM · TEAM LEAD
>
> *"That could also help with understanding how to prompt better, because most people don't know."*
>
> PMM · ECOSYSTEM AND PLATFORM

7.2.1 Validation outcome

PMMs reviewed the guide in a workshop and confirmed that both the prompts and the staged structure were useful for their work — judging the structure a way to organise the work rather than an extra layer on top of it. Against the three SQ3 criteria: **relevance is high** (the encoded pain points are the ones PMMs raised unprompted); **desirability is supported** by their confirmation that prompts and structure are useful in practice; and **feasibility is grounded** in delivering through Copilot inside the governed M365 environment they already work in, so adoption requires no new platform.

TRANSLATION LAYER

The RMM sits between central marketing and local market reality — adapting input down, feeding intelligence back up.

"We partially have too much information. And we don't have enough resources to structure this information."

RMM · CENTRAL REGION

7.3 Secondary validation: Regional Market Managers

The two RMM interviews describe the same structural position from two markets: a translation layer taking central strategy, campaigns, and messaging and adapting them into activities that work in the region, while feeding local intelligence back upward.

> *"Translate the central input that I receive into actionable marketing activities in the field."*
>
> RMM · SOUTHERN EUROPE

Three problems recur: annual-planning overload when many sources arrive at once with no analytical support; continuous, manual competitive monitoring that is never finished; and a localisation gap across language, channel, and regulation. The **synthetic RMM persona** composites these into a regional manager on a small team, reseller-heavy route to market, analytical load concentrated in the planning cycle, willing but not expert at prompting — one RMM explicitly asked for centralised prompting support.

7.3.1 Anticipated fit against the guide

DOCUMENTED RMM NEED	STAGE	ANTICIPATED FIT
Competitive monitoring across verticals	Prioritise · Reduce	Prompts terminate in a sales-ready positioning response, converting monitoring into reseller-facing output rather than another summary.
Annual-planning overload, no analytical support	Reduce · Interpret	The agreement-and-conflict prompt structure forces triangulation across sources, addressing the structuring gap directly.
Localisation across language, channel, regulation	Mobilise	Prompts name the adaptation axis and compress central content for reseller use, leaving cultural validation to the human.

Relevance is strong and **feasibility favourable** (template-based, low prompting skill, runs in Copilot); desirability is supported but inferred — the honest limit is anticipated fit grounded in real needs, not confirmed adoption.

DIRECT EVIDENCE

A written response describes the guide's intended behaviour without having seen it — the strongest desirability signal for Sales.

"The challenge today is not a lack of information, but an abundance of fragmented information."

SALES · WRITTEN RESPONSE

7.4 Secondary validation: Sales

Sales evidence comes from three participants — two interviews and one detailed written response. Insight is used most heavily at the front of the sales cycle, drawing chiefly on PMM output. Sales overload is fragmentation, not volume, more clearly here than in any other function: pre-call research that consumes time per lead, difficulty finding the right internal document, and the conversion of long product updates into usable knowledge.

> *"An hour of information is a lot to process … they cannot remember a one-hour webinar and really use that in their conversation."*
>
> SALES · INSIDE SALES LEAD

7.4.1 Direct evidence: the written response

One participant's structured written response functions as direct empirical evidence of desirability, describing — unprompted, without sight of the guide — almost exactly what it is built to deliver: a customer name turned into a structured briefing, trust conditioned on updated, verifiable, source-checkable information.

> *"The best information for a sales team is not the longest, but the one that enables fast and confident action."*
>
> SALES · WRITTEN RESPONSE

DOCUMENTED SALES NEED	STAGE	ANTICIPATED FIT
Repeated, time-bound pre-call research	Prioritise · Reduce	A fill-in-the-blank prompt produces a structured pre-call briefing, matching the request for a name-to-briefing tool.
Finding the right document among underused content	Reduce	A prompt interrogates a single document for a specific customer question, addressing discoverability at the point of use.
Converting hour-long updates into usable knowledge	Reduce · Interpret	Prompts compress webinar and update content into retained, conversation-ready points.

The Sales assessment is the strongest of the two secondary groups: **relevance high and triangulated** across three participants, **desirability with direct empirical support**, feasibility favourable.

ONE LINE

Across all three functions the problem is fragmentation and the burden of trusting information — never its quantity.

7.5 Cross-role validation synthesis

Read together, the three functions validate the guide along a consistent line. In every group the problem is fragmentation and the burden of trusting information, not its quantity: PMMs locate the cost in *validating* data, RMMs in *structuring and translating* simultaneous inputs, Sales in *compressing and retrieving* content for conversation. The guide answers each with the same underlying move — a staged reduction from attention to action, expressed as concrete Copilot prompts underpinned by a shared trust-and-verification context.

The evidence is strongest where it should be. PMMs, the primary stakeholders, validated the finished artefact directly. Sales has independent written confirmation of desirability. The RMM case rests on close alignment between documented needs and the relevant prompts — convincing on relevance, lighter on confirmed desirability. Across all three, the trust-and-verification design responds to the conditional, hallucination-scaled trust that recurred throughout the research, which is what gives the guide a credible claim to being *adopted* rather than merely understood.

7.6 Limitations and next validation steps

The validation has clear boundaries. For the secondary stakeholders, anticipated fit is not confirmed adoption; the Sales desirability evidence, though strong, rests on a single written response; and the personas are composites from a small number of interviews — analytical instruments, not representative samples. These boundaries define the immediate next steps, aligned with the first horizon of the implementation roadmap: a short live walkthrough of the RMM and Sales sections with two or three practitioners each to convert anticipated fit into observed use; a small pilot of the pre-call, localisation, and annual-planning prompts under genuine time pressure; and the shared prompt-library governance of the Equip layer to maintain and improve validated prompts over time. None of these changes the conclusion of this chapter; they deepen the evidence beneath it.

CHAPTER SUMMARY → Validation confirms the guide is **relevant, desirable, and feasible**, with empirical confirmation from the primary PMM group, direct written evidence from Sales, and strong anticipated fit for RMMs — all converging on fragmentation and trust as the shared problem the guide resolves. With the artefact validated, Chapter 8 presents the final deliverables in full: the PRISM(E) Prompt Guide, the strategic and tactical roadmaps, and the Equip enablement components that carry the framework into sustained practice.

CHAPTER EIGHT

08 Final Deliverables

Three connected artefacts: a prompt guide, a strategic roadmap, and a tactical roadmap.

The project delivers three connected artefacts. The AI Prompt Guide makes the framework usable at the individual level, while the strategic and tactical roadmaps describe how Webfleet can move PRISM(E) from a personal practice to an organisational capability.

8.1 The PRISM(E) Prompt Guide

The Prompt Guide is the working output at the individual level: a practical toolkit for reducing overload and improving GtM insight synthesis using Microsoft Copilot across three roles. It opens with a short how-to and a one-page view of the six stages, then provides prompts for each stage in the three-part form of a copy-paste prompt, a template prompt, and a statement of prompt logic. A solution map translates the PMM pain points from Chapter 5 into specific prompt approaches; dedicated sections then provide role-specific prompts for PMMs, RMMs, and Sales. The guide closes with the Equip section and a reusable context prompt. Its governing principle is stated throughout: AI is a synthesis and reduction tool, not an autonomous decision-maker.

8.1The Prompt Guide in practice

Four representative spreads show how the guide turns PRISM(E) into day-to-day practice: a one-page view of the six stages and the three-part prompt form, a worked stage page, and two reference tools that connect the prompts back to the research. The full twenty-eight-page guide is reproduced in Appendix D.



HOW TO USE · SIX-STAGE VIEW

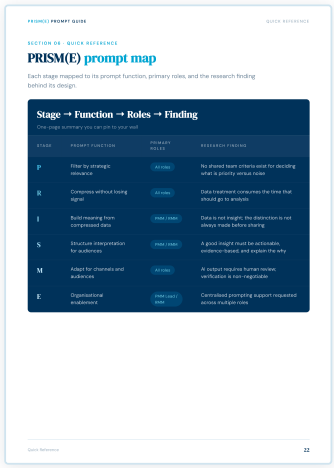


A STAGE PAGE · THREE-PART PROMPT FORM

Fig. 8.1 The guide's structure — a one-page view of the six PRISM(E) stages with the three-part prompt form (left), and a worked stage page giving a copy-paste prompt, a fill-in template, and the prompt logic (right).



SOLUTION MAP · PAIN POINT → PROMPT



QUICK REFERENCE · STAGE → FINDING

Fig. 8.2 Two reference tools — the solution map translating the Chapter 5 PMM pain points into specific prompt approaches (left), and the one-page quick-reference linking each stage to its prompt function, primary roles, and the research finding behind it (right).

18 MONTHS

A future vision of a Product Marketing function that thinks with AI rather than around it.

8.2 Strategic roadmap: from pilot to permanent capability

The strategic roadmap sets out an eighteen-month, three-horizon plan for embedding PRISM(E) into Product Marketing, owned by the PMM Lead and running on Microsoft Copilot (Fig. 8.3).

HORIZON 1
H1
Months 1–4

Foundation & Pilot

Establishes the infrastructure and validates PRISM(E) in controlled conditions with a small PMM cohort, producing the first AI-assisted GtM briefs.

HORIZON 2
H2
Months 5–10

Scale & Integration

Rolls the framework across PMMs and RMMs, integrates it with the quarterly GtM cycle, and formalises the Synthesise-to-Mobilise handoff.

HORIZON 3
H3
Months 11–18

Institutionalise & Evolve

Embeds PRISM(E) as a permanent capability, measures longer-term impact on GtM quality, and iterates the framework on lived practice.

The Equip layer runs across all three horizons, because organisational enablement, AI governance, GDPR compliance, leadership sponsorship, and capability-building are ongoing conditions rather than one-time activities. Without a healthy Equip layer, the five sequential stages degrade.

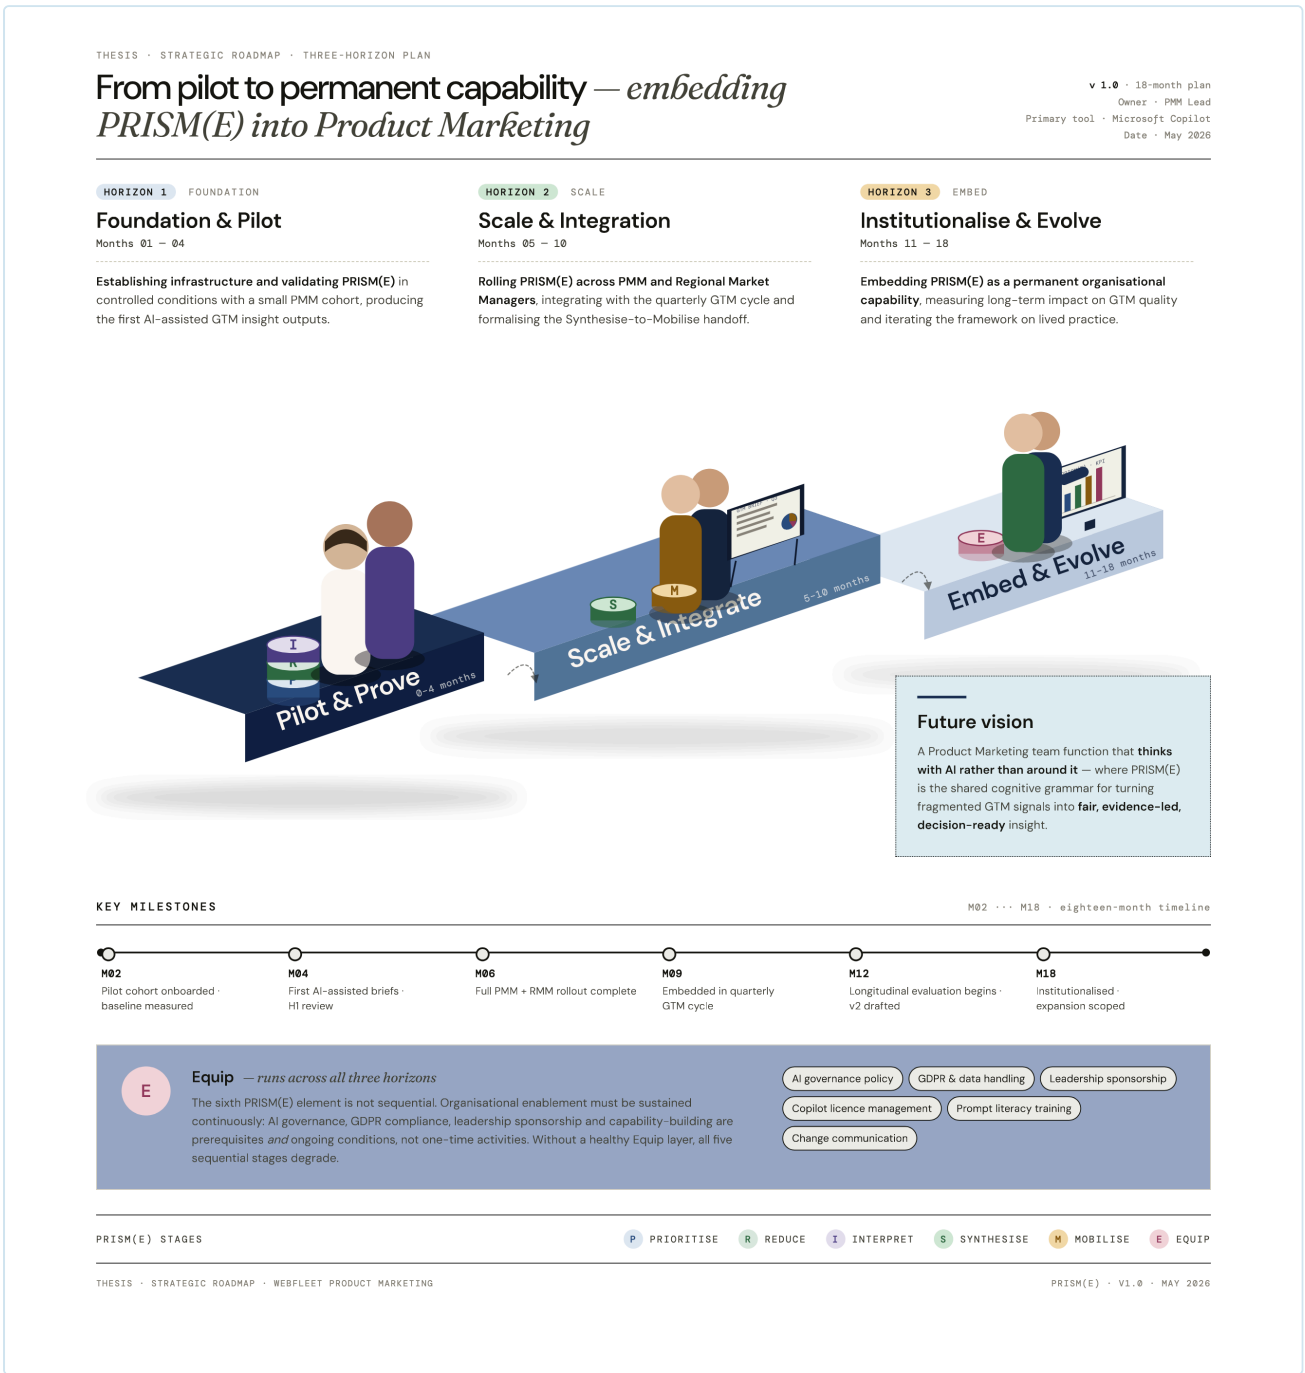


Fig. 8.3 Strategic roadmap — from pilot to permanent capability: the eighteen-month, three-horizon plan for embedding PRISM(E) into Product Marketing, with the Equip layer running across all three horizons.

FIVE LENSES

Adoption planned as organisational change, not a tool rollout.

8.3 Tactical roadmap: building an AI-ready practice

The tactical roadmap reads the same eighteen-month journey through five lenses, so adoption is planned as an organisational change rather than a tool rollout (Fig. 8.4).

Direction & context
From an AI-governance and GDPR baseline to a searchable knowledge base supported by agents.

People & roles
PMMs as pilots → mentors → co-authors; RMMs as observers → co-creators; Sales as receivers; leadership as stewards.

Ways of working
From filtering with intent to synthesising side-by-side from one structured brief guarded by the audit gate.

Related values
Trust, clarity, workability, shared ownership, verification, collective sensemaking, sustainable practice.

A fifth lens — **outcomes & risks** — tracks the closing of the four gaps and, importantly, surfaces shadow AI use so that informal prompts and workarounds are pulled into the shared library and inform the next version rather than remaining hidden.

8.4 The Equip enablement components

Across both roadmaps, the Equip layer is operationalised through three components: a **shared prompt library** with governance and a quarterly review owned by the PMM Lead; **tool-specific onboarding** that names known weaknesses and verification triggers; and an **organisational access strategy** so outputs can be shared without tool-access gaps. These are the difference between PRISM(E) as a personal habit and PRISM(E) as a team capability.

CHAPTER SUMMARY → The project delivers three connected artefacts — the prompt guide at the individual level, and strategic and tactical roadmaps with Equip components at the organisational level — that together carry PRISM(E) from a personal practice into a sustained team capability. Chapter 9 discusses the theoretical contribution and practical implications of the work.

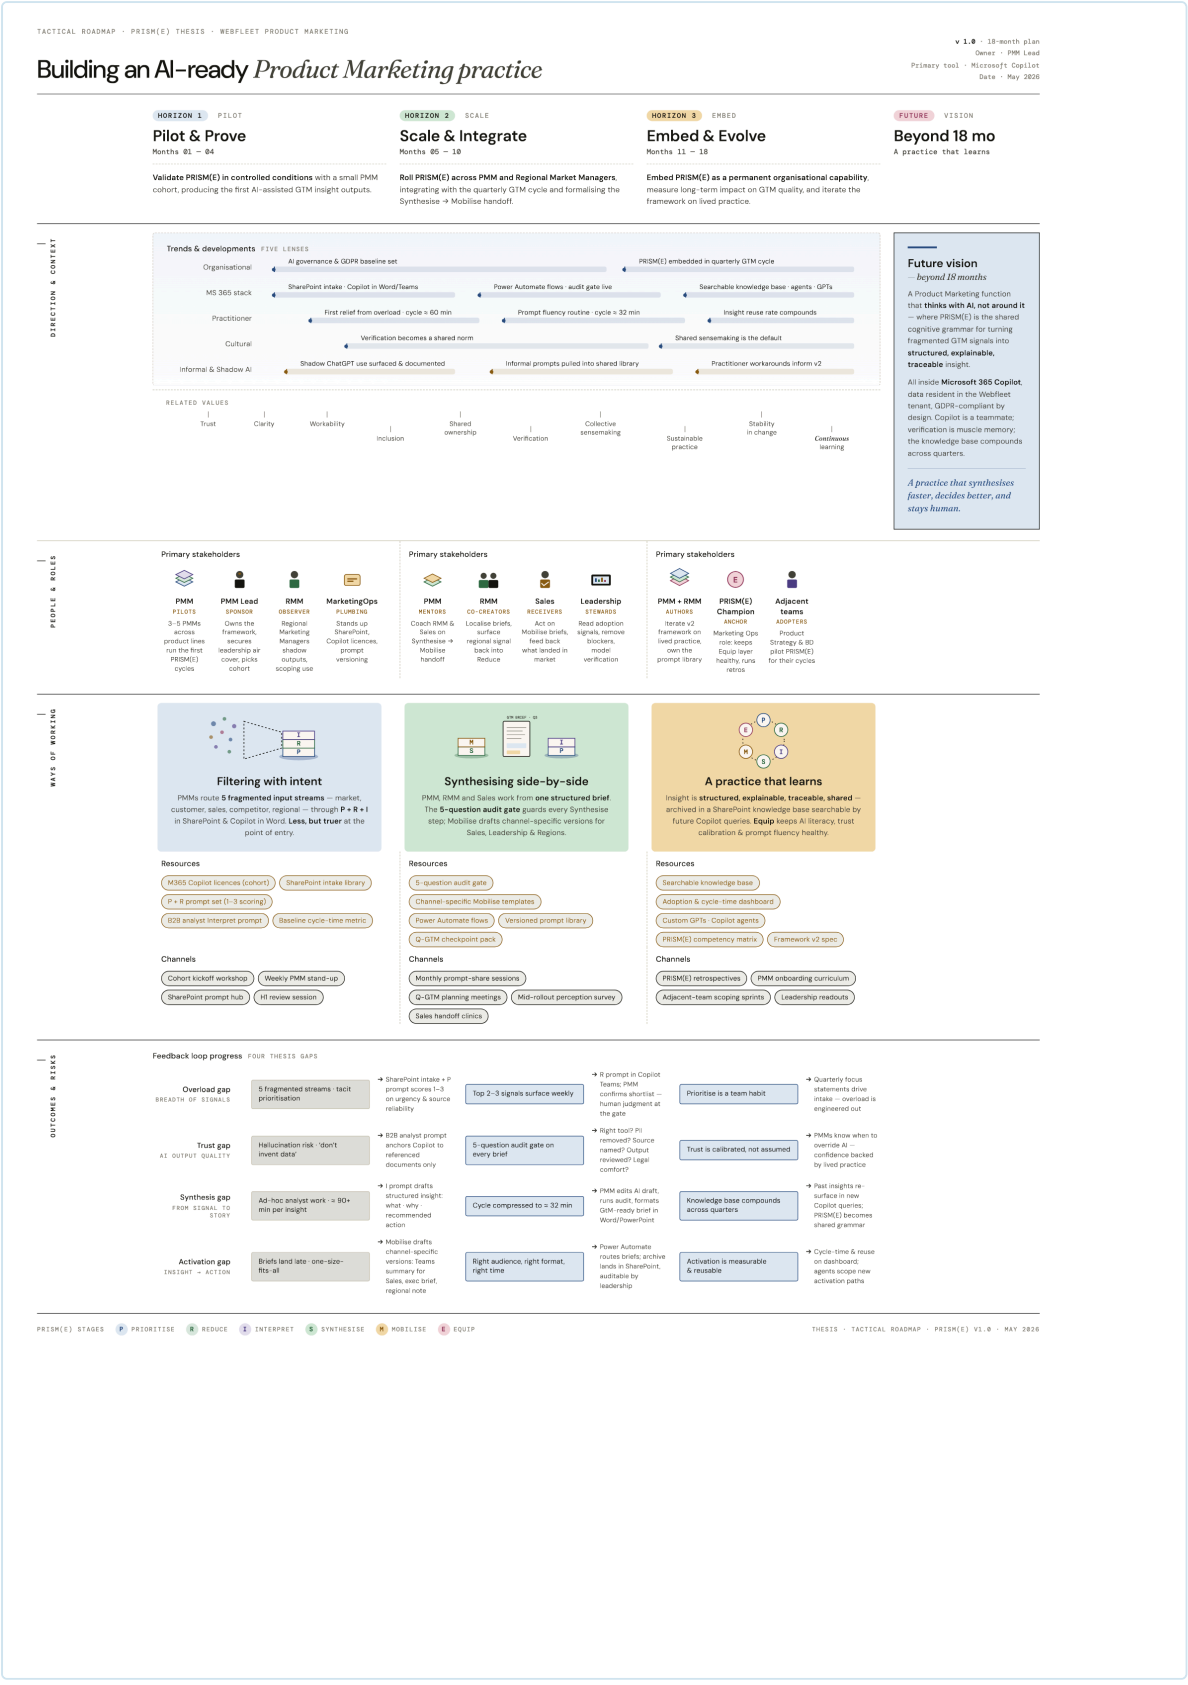


Fig. 8.4 Tactical roadmap — building an AI-ready practice: the same eighteen-month journey read through five lenses (direction & context, people & roles, ways of working, related values, and outcomes & risks).

CHAPTER NINE

09 Discussion and Conclusion

9.1 Answering the research question

The project asked how a structured, AI-enabled sensemaking framework can reduce information overload and improve the quality of GtM insight synthesis for PMMs at Webfleet. The answer developed across the three sub-questions.

SQ1 Established **where overload arises.** PMMs do not lack information; they lack shared criteria for filtering it, lose analytical time to data treatment, and often share data as if it were insight. Overload is experienced as breadth, depth, and fragmentation — so any intervention must be structured enough to be consistent yet flexible enough to fit different tasks.

SQ2 Produced the **design response.** PRISM(E) structures the workflow into Prioritise, Reduce, Interpret, Synthesise, and Mobilise, with an Equip layer — preserving human judgment at every consequential gate rather than automating it away. Operationalising it inside Copilot within Webfleet's own tenant grounds it in the organisation's real environment and governance constraints.

SQ3 Addressed through **perception-based assessment**, completed in the validation chapter. On the evidence to date, the framework's structure reflects how practitioners already reason, and the co-creation sessions indicate the prompt library is recognisable as a response to real tasks rather than an imposition.

CONTRIBUTION

Five streams, one role: the contribution is the connective synthesis, not any single stage.

9.2 Theoretical contribution

The central theoretical contribution is the synthesis of five previously separate streams into a single, empirically grounded workflow for a role none of them has addressed directly: the go-to-market insight-synthesis work of B2B Product Marketing Managers. Positioned against the literature reviewed in Chapter 2, the contribution can be stated stream by stream, in each case as an extension beyond where that literature currently stops.

Information overload and technostress. Overload has been theorised as multi-factorial but located in managerial, accounting, and MIS settings (Eppler et al., 2004), while the technostress literature names techno-overload as a stressor without examining generative AI (Ragu-Nathan et al., 2008; Tarafdar et al., 2007). This study extends both to the PMM role and shows AI occupying a dual position — simultaneously a potential cause and a remedy of overload — that neither literature articulates.

Sensemaking. Sensemaking is widely cited but rarely operationalised as design (Weick, 1995). This research converts its properties of bracketing and enactment into explicit architectural principles for an AI workflow, and extends prior academic-knowledge-work research into a commercial go-to-market context (Hao et al., 2025).

Prompt engineering. Prompt patterns have been established in software, educational, and general-productivity settings (White et al., 2023; Reynolds et al., 2021); this study contributes a role-specific, organisation-grounded application — a prompt library tied to one team's tasks, data, and governance rather than to generic productivity.

B2B knowledge creation and innovation management. AI value has been located in the structured process rather than the technology (Fosso Wamba et al., 2021; Mariani et al., 2023); this research specifies what that process looks like for product marketing and shows how it compounds through a shared, reusable insight archive.

Situated AI and adoption. Situated-AI theory has been supplied at a strategic level (Enholm et al., 2023), and workforce readiness identified as decisive but with recommendations that stop at training (Kumar et al., 2025; Ghaly et al., 2024). This study answers that open call with a concrete design response — the Equip layer — that operationalises grounding, bounding, and recasting inside a functional workflow.

Taken together, the contribution is less any single stage than the connective synthesis itself: an evidence-led demonstration that overload, sensemaking, prompting, knowledge creation, and situated AI can be integrated into one coherent, role-specific design. PRISM(E) is the artefact through which that integration is made visible and testable.

ACTION

Buying AI capability does not produce adoption \u2014 the workflow and the readiness layer have to be designed together.

9.3 Practical implications

The findings translate into concrete actions, framed first for Webfleet and then for comparable organisations. For Webfleet, the framework makes insight synthesis consistent, auditable, and reusable without removing the human judgment that makes an insight defensible. Seven actions follow directly from the evidence.

RECOMMENDED ACTIONS FOR WEBFLEET

1. **Run a bounded pilot.** Trial PRISM(E) for one quarterly GtM cycle with two or three PMMs on a live priority, inside the existing M365 tenant \u2014 no procurement or new tooling required.
2. **Make the intake and archive real now.** Stand up the shared SharePoint intake folder and insight library, and adopt the numerical-anchor rule and five-question Synthesise audit gate as the team standard for any externally shared insight.
3. **Centralise prompting.** Publish the prompt library and appoint a prompt owner per team \u2014 the direct answer to the centralised-support request from RMM2 and Sales.
4. **Close the discoverability gap.** Tag and index PMM material so Copilot can retrieve it, and give Sales the fill-in-the-blank prospect-brief prompts they asked for \u2014 reclaiming the ~20 minutes per lead Sales1 loses to manual research.
5. **Pull PMMs upstream.** Involve product marketing earlier in projects to relieve the late-involvement pressure no tool can fully resolve.
6. **Invest in Equip before scaling.** Run AI-literacy sessions, declare each tool's known weaknesses, and tier trust by stakes; readiness is the precondition for adoption, not an afterthought.
7. **Investigate forecast reliability.** Examine the team-wide sales-forecast over-promising directly with the Sales function, since it is a structural source-reliability issue rather than an individual one.

For other organisations, the transferable lesson is twofold. First, the approach generalises to any B2B product-marketing function working with fragmented inputs inside a Microsoft 365 environment; the stages are tool-agnostic in principle. Second, and more importantly, the case shows that buying AI capability does not by itself produce adoption \u2014 value comes from designing the workflow and the readiness layer together, keeping a human gate at every consequential step. Organisations that treat AI as a filter-and-augment discipline rather than an automation shortcut are the ones positioned to turn it into durable advantage.

REFLECTION

"Let the framework follow the evidence rather than the other way around."

9.4 Limitations

Several limitations bound the claims of this work. The study is situated in a single organisation, which strengthens contextual fit but limits generalisability. The sample is small and qualitative, appropriate to the research questions but not a basis for statistical claims. The framework testing could not be lived-in within the thesis window, so assessment is perception-based rather than a measure of realised effectiveness. The refinement from v1 to v2 was largely additive; a longer study would test the framework against disconfirming evidence and sustained use.

9.5 Recommendations and future research

Future work should follow the framework into lived practice over at least one full quarterly GtM cycle, measuring not only perceived usefulness but realised changes in cycle time, insight reuse, and decision quality. The sales-forecast reliability finding deserves direct study with the sales function itself. The expansion of PRISM(E) to adjacent functions, and the development of the specialised agents practitioners proposed, are natural next steps once the framework is embedded.

9.6 Personal reflection

This project sat at the boundary between research and practice that drew me to Strategic Product Design. The most useful discipline I learned was to let the framework follow the evidence rather than the other way around. The harder and more rewarding work was to keep returning to what practitioners actually said — including the parts that complicated the design. Working inside a real organisation taught me that a framework is only as good as its fit with how people already think, which is why the Equip layer and the roadmaps came to feel as important as the five sequential stages. I leave the project more convinced that AI is most valuable in knowledge work not when it replaces judgment but when it gives judgment better material to work with.

IN CLOSING The study answers its research question by showing that a structured, AI-enabled sensemaking workflow reduces overload when it filters before it interprets and preserves a human gate at every consequential step. Its lasting contribution is the connective synthesis of five literature streams into one role-specific, validated practice — and a concrete path for Webfleet to adopt it.

TRANSPARENCY

Generative AI was used as a thinking and drafting aid throughout this project; all judgments, claims, and conclusions remain the author's own.

USE OF AI

Use of AI

Because this thesis is itself a study of AI in knowledge work, generative AI tools were used deliberately and reflectively during the research and design process, both as objects of study and as working aids. This section documents which tools were used and how, in the interest of transparency.

Tools used

Microsoft Copilot

Used within the Microsoft 365 environment to explore how an enterprise-grade, organisationally embedded assistant behaves in practice — directly informing the design of the PRISM(E) stages and the Equip enablement layer.

ChatGPT (OpenAI)

Used to test and compare model capabilities, to brainstorm and pressure-test early ideas, and to help structure and sharpen written sections of the report.

Claude (Anthropic)

Used for extended reasoning and drafting — organising arguments, refining the framing of the framework and deliverables, and stress-testing the structure of the thesis.

How these tools were used

The tools served four purposes across the project. First, **capability testing**: the assistants were used hands-on to understand what current AI tools can and cannot do for go-to-market sensemaking — evidence that fed directly into the framework. Second, **ideation**: they supported brainstorming and the rapid exploration of alternative directions during concept development. Third, **writing and structure**: they helped organise arguments, tighten prose, and improve the readability of the report. Fourth, **framing the deliverables**: they assisted in shaping and articulating the framework, prompt library, and roadmaps.

In every case, the AI was used as a support to thinking, not a substitute for it. All sources, analysis, interpretation of findings, and design decisions were reviewed, verified, and authored by the researcher — consistent with the central argument of this thesis: that “AI is most valuable when it gives human judgment better material to work with, not when it replaces that judgment”.

REFERENCES

Selected sources across the five theoretical streams.

BIBLIOGRAPHY

References

ABI Research. (2025, February 20). *Geotab, Samsara, and PowerFleet take the top spots in ABI Research's commercial telematics competitive ranking* [Press release]. PR Newswire. prnewswire.com/news-releases/geotab-samsara-and-powerfleet-take-the-top-spots-in-abi-researchs-commercial-telematics-competitive-ranking-302380732.html

Adadi, A., & Berrada, M. (2018). Peeking inside the black box: A survey on explainable artificial intelligence (XAI). *IEEE Access, 6,* 52138–52160. doi.org/10.1109/ACCESS.2018.2870052

Agrawal, R., et al. (2025). Applications of generative AI in business analytics: Transforming decision-making and strategy. *IEEE.* doi.org/10.1109/InC465408.2025.11256227

Bridgestone. (2024). *Webfleet and Bridgestone join forces for full mobility solutions.*

Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research. doi.org/10.3386/w31161

Cano-Marín, E., et al. (2025). Market research and knowledge using generative AI: The power of large language models. *Journal of Innovation & Knowledge.* doi.org/10.1016/j.jik.2025.100796

Dell'Acqua, F., et al. (2023). *Generative AI at work.* National Bureau of Economic Research. doi.org/10.3386/w31161

Dwivedi, Y. K., et al. (2025). Artificial intelligence adoption in marketing strategies. *Social Sciences & Humanities Open.* doi.org/10.1016/j.ssaho.2025.102048

Enholm, I. M., Papagiannidis, E., Mikalef, P., & Krogstie, J. (2023). Competitive advantage through artificial intelligence: Toward a theory of situated AI. *Academy of Management Review.* doi.org/10.5465/amr.2020.0205

Eppler, M. J., & Mengis, J. (2004). The concept of information overload: A review of literature from organization science, accounting, marketing, MIS, and related disciplines. *The Information Society, 20*(5), 325–344. doi.org/10.1080/01972240490507974

Fosso Wamba, S., Queiroz, M. M., & Trinchera, L. (2021). An integrated artificial intelligence framework for knowledge creation and B2B marketing rational decision making for improving firm performance. *Industrial Marketing Management, 92,* 178–189. doi.org/10.1016/j.indmarman.2020.12.001

Frost & Sullivan. (2025). *Frost & Sullivan names Webfleet European fleet telematics company of the year 2025.* Bridgestone Press.

Ghaly, M. A., Mubarak, M. F., & Eldesouky, A. (2024). Digital transformation potential: The role of artificial intelligence in business. *Business Review, 9*(3). doi.org/10.26668/businessreview/2024.v9i3.4499

GlobeNewswire. (2026, January 12). *Europe fleet management industry report 2025–2029.*

Hao, Y., Zhang, Y., et al. (2025). The design space of recent AI-assisted research tools for ideation, sensemaking, and scientific creativity. *arXiv.* doi.org/10.48550/arXiv.2502.16291

Kumar, V., et al. (2025). Artificial intelligence (AI) adoption in marketing strategies: Navigating the present and shaping the future business landscape. *Social Sciences & Humanities Open.* doi.org/10.1016/j.ssaho.2025.102048

Li, X., et al. (2025). From emotion to reflection: Leveraging EmotionPrompt strategy to empower self-determination in decision-making with generative artificial intelligence. *Information & Management.* doi.org/10.1016/j.im.2025.104194

Lund, B., et al. (2024). A literature review of artificial intelligence research in business and management using machine learning and ChatGPT. *Data and Information Management.* doi.org/10.1016/j.dim.2024.100076

Mariani, M. M., Machado, I., Nambisan, S., & Dwivedi, Y. K. (2023). Artificial intelligence in innovation management: A review of innovation capabilities and a taxonomy of AI applications. *Journal of Product Innovation Management.* doi.org/10.1111/jpim.12698

Nwafor, C., et al. (2024). Explainable AI: The quest for transparency in business and beyond. *IEEE.* doi.org/10.1109/ICICT62343.2024.00093

Paschen, J., et al. (2025). Market research using generative AI. *Journal of Innovation & Knowledge.* doi.org/10.1016/j.jik.2025.100796

Peeters, M., Klievink, B., et al. (2024). Generative artificial intelligence in innovation management: A preview of future research developments. *Journal of Business Research.* doi.org/10.1016/j.jbusres.2024.114542

Ragu-Nathan, T. S., Tarafdar, M., Ragu-Nathan, B. S., & Tu, Q. (2008). The consequences of technostress for end users in organizations: Conceptual development and empirical validation. *Information Systems Research, 19*(4), 417–433. doi.org/10.1287/isre.1070.0165

Reynolds, L., & McDonell, K. (2021). Prompt programming for large language models: Beyond the few-shot paradigm. *arXiv.* doi.org/10.48550/arXiv.2102.07350

Tarafdar, M., Tu, Q., Ragu-Nathan, B. S., & Ragu-Nathan, T. S. (2007). The impact of technostress on role stress and productivity. *Journal of Management Information Systems, 24*(1), 301–328. doi.org/10.2753/MIS0742-1222240109

Webfleet. (2024). *About Webfleet: Company history, mission and vision.*

Weick, K. E. (1995). *Sensemaking in organizations.* SAGE Publications.

White, J., Fu, Q., Hays, S., Sandborn, M., Olea, C., Gilbert, H., Elnashar, A., Spencer-Smith, J., & Schmidt, D. C. (2023). A prompt pattern catalog to enhance prompt engineering with ChatGPT. *arXiv.* doi.org/10.48550/arXiv.2302.11382

Ye, R., et al. (2025). AI-assisted sensemaking tools. *arXiv.* doi.org/10.48550/arXiv.2502.16291

A–G

SUPPORTING MATERIAL

Appendices

A	Interview questions guide
B	Informed consent form
C	Data management plan
D	The PRISM(E) Prompt Guide
E	Co-creation workshop materials (Miro)
F	HREC
G	Strategic and tactical roadmaps

APPENDIX A

Interview questions guide

Semi-structured interview guide for the contextual inquiry — six identical questions across all three stakeholder groups, each with one or two role-specific probes. 45 minutes per interview: Opening (15) · Core (15) · Closing (15).

ITEM	DETAIL
Duration	45 minutes per interview.
Structure	Opening (15 min) · Core (15 min) · Closing (15 min).
Questions	Six core questions, identical across all groups.
Probes	One or two role-specific probes per question, adapted to each group's context.
Groups	Product Marketing Managers (primary) · Sales team · Regional market leads (secondary).
Recording	Audio recorded with verbal consent · transcribed for thematic analysis.
Analysis	Theoretically directed thematic analysis.
Ethics	Participation voluntary · data anonymised in thesis reporting.

The six core questions

#	THEME	QUESTION (IDENTICAL ACROSS ALL GROUPS)	GENERIC PROBE
Q1	**Role & workflow**	**Can you brief me about your role and how GtM insights fit into your day-to-day work?**	*What does a typical week look like for you?*
Q2	**Sources & volume**	**How many sources of information do you typically work with, and which feel most useful versus most overwhelming?**	*Has the volume of information you receive changed recently?*
Q3	**Overload experience**	**Explain the last time you had to make sense of a large amount of information — where did it feel most overwhelming or uncertain?**	*What did you do when that happened?*
Q4	**Relevance filtering**	**How do you currently decide what information is relevant enough to act on versus what to ignore or set aside?**	*Is there a process or rule you follow, or is it more instinctive?*
Q5	**Insight quality**	**What does a good insight look like to you, and what makes one fall short?**	*Can you give me a recent example of each?*
Q6	**AI readiness & trust**	**Do you currently use any AI tools in your work, and what would it take for you to trust an AI-assisted insight enough to act on it?**	*What would make you not trust it?*

APPENDIX A · CONTINUED

Role-specific probes

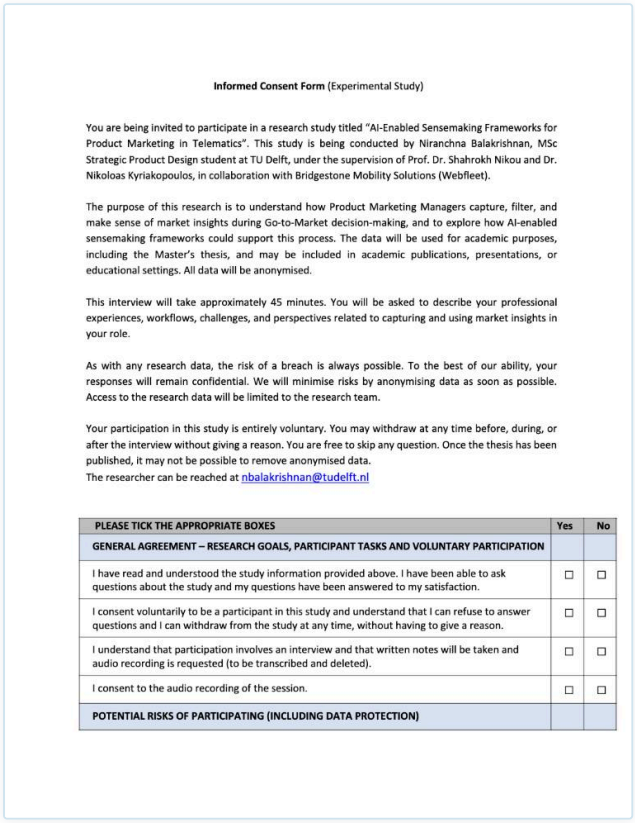
The same six questions were asked of every group; the probes below were tailored to each role's vantage point on the GtM chain.

#	THEME	PRODUCT MARKETING	SALES TEAM	REGIONAL LEADS
Q1	**Role & workflow**	How much of your week is spent processing or synthesising market information?	At what point in your sales process do you typically use GtM insights from PMMs?	What does the flow of information between central and your region look like in a typical week?
Q2	**Sources & volume**	Which sources feel most reliable for GtM decisions, and which do you find yourself doubting?	Beyond what PMMs send you, what other sources do you rely on before customer conversations?	How much of your information comes from central versus local or regional sources?
Q3	**Overload experience**	Did you feel you had enough time to process everything before the decision had to be made?	Recall a moment with a lot of information before a customer conversation — where were you most uncertain about what to use?	Was the overload coming from too much central information, too much local complexity, or both?
Q4	**Relevance filtering**	Do you have shared criteria with your team for deciding what is priority versus noise?	When a batch of insights arrives from the PMM team, what is the first thing you look at to judge relevance for your accounts?	When central insights arrive, do you apply them as-is or always translate them first — and what drives that?
Q5	**Insight quality**	How do you know when an insight is ready to share versus when it still needs work?	What would a perfect PMM insight brief look like before a key customer meeting?	—
Q6	**AI readiness & trust**	If AI summarised your weekly intake, would you feel confident presenting it in a GtM meeting — why or why not?	What would reassure you the AI output was grounded in real, current data rather than something generic?	If AI adapted central insights for your region, what would you check before trusting that adaptation?

APPENDIX B

Informed consent form

The opt-in consent form signed by every interview participant before recording, covering purpose, voluntary participation, data handling, and anonymisation.



Informed Consent Form (Experimental Study)

You are being invited to participate in a research study titled "AI-Enabled Sensemaking Frameworks for Product Marketing in Telematics". This study is being conducted by Niranchna Balakrishnan, MSc Strategic Product Design student at TU Delft, under the supervision of Prof. Dr. Shahrokh Nikou and Dr. Nikoloas Kyriakopoulos, in collaboration with Bridgestone Mobility Solutions (Webfleet).

The purpose of this research is to understand how Product Marketing Managers capture, filter, and make sense of market insights during Go-to-Market decision-making, and to explore how AI-enabled sensemaking frameworks could support this process. The data will be used for academic purposes, including the Master's thesis, and may be included in academic publications, presentations, or educational settings. All data will be anonymised.

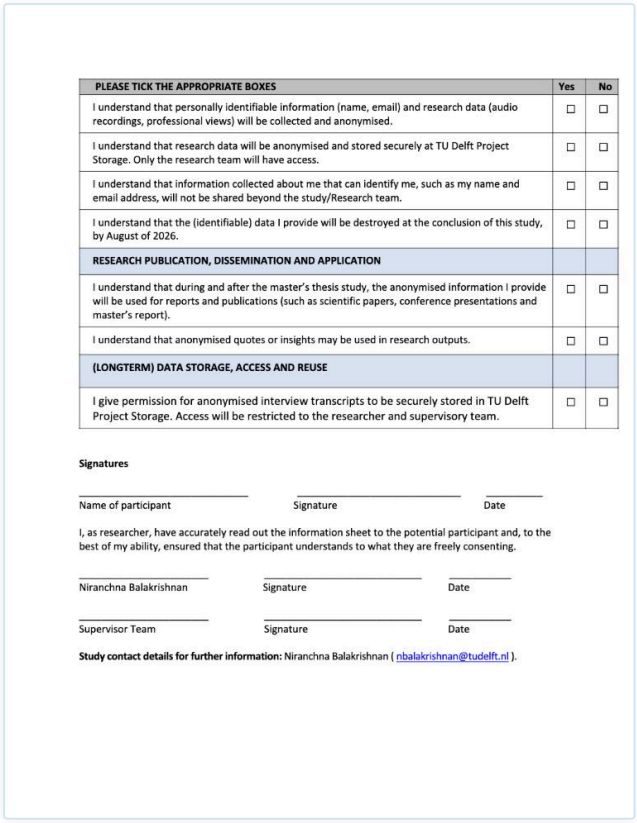
This interview will take approximately 45 minutes. You will be asked to describe your professional experiences, workflows, challenges, and perspectives related to capturing and using market insights in your role.

As with any research data, the risk of a breach is always possible. To the best of our ability, your responses will remain confidential. We will minimise risks by anonymising data as soon as possible. Access to the research data will be limited to the research team.

Your participation in this study is entirely voluntary. You may withdraw at any time before, during, or after the interview without giving a reason. You are free to skip any question. Once the thesis has been published, it may not be possible to remove anonymised data.
The researcher can be reached at nbalakrishnan@tudelft.nl

PLEASE TICK THE APPROPRIATE BOXES	Yes	No
GENERAL AGREEMENT – RESEARCH GOALS, PARTICIPANT TASKS AND VOLUNTARY PARTICIPATION		
I have read and understood the study information provided above. I have been able to ask questions about the study and my questions have been answered to my satisfaction.	☐	☐
I consent voluntarily to be a participant in this study and understand that I can refuse to answer questions and I can withdraw from the study at any time, without having to give a reason.	☐	☐
I understand that participation involves an interview and that written notes will be taken and audio recording is requested (to be transcribed and deleted).	☐	☐
I consent to the audio recording of the session.	☐	☐
POTENTIAL RISKS OF PARTICIPATING (INCLUDING DATA PROTECTION)		

PAGE 1 · INFORMATION & CONSENT



PLEASE TICK THE APPROPRIATE BOXES	Yes	No
I understand that personally identifiable information (name, email) and research data (audio recordings, professional views) will be collected and anonymised.	☐	☐
I understand that research data will be anonymised and stored securely at TU Delft Project Storage. Only the research team will have access.	☐	☐
I understand that information collected about me that can identify me, such as my name and email address, will not be shared beyond the study/Research team.	☐	☐
I understand that the (identifiable) data I provide will be destroyed at the conclusion of this study, by August of 2026.	☐	☐
RESEARCH PUBLICATION, DISSEMINATION AND APPLICATION		
I understand that during and after the master's thesis study, the anonymised information I provide will be used for reports and publications (such as scientific papers, conference presentations and master's report).	☐	☐
I understand that anonymised quotes or insights may be used in research outputs.	☐	☐
(LONGTERM) DATA STORAGE, ACCESS AND REUSE		
I give permission for anonymised interview transcripts to be securely stored in TU Delft Project Storage. Access will be restricted to the researcher and supervisory team.	☐	☐

Signatures

Name of participant _____ Signature _____ Date _____

I, as researcher, have accurately read out the information sheet to the potential participant and, to the best of my ability, ensured that the participant understands to what they are freely consenting.

Niranchna Balakrishnan _____ Signature _____ Date _____

Supervisor Team _____ Signature _____ Date _____

Study contact details for further information: Niranchna Balakrishnan (nbalakrishnan@tudelft.nl).

PAGE 2 · SIGNATURES

Appendix B Informed consent form (two pages, reproduced at reduced size). The signed originals are held in the project's data-management record.

APPENDIX C

Data management plan

The TU Delft data management plan (DMP) for the study, covering data collection, storage, anonymisation, retention, and sharing, reproduced at reduced size (8 pages).



Plan Overview

A Data Management Plan created using DMPonline

Title: AI-Enabled Sensemaking Frameworks for Product Marketing in Telematics

Creator: Niranchna Balakrishnan

Affiliation: Delft University of Technology

Template: TU Delft Data Management Plan template (2025)

Project abstract:
This research explores how Product Marketing Managers (PMMs) in the telematics industry capture, filter, and make sense of fragmented market insights during Go-to-Market decision-making. Using a qualitative, exploratory approach grounded in Sensemaking theory, the study is conducted in collaboration with Bridgestone Mobility Solutions (Webfleet) and involves qualitative interviews and feedback sessions with approximately 8–15 participants in professional roles. The research aims to identify current challenges in insight handling and develop a decision-tree-based sensemaking framework and supporting AI-assisted workflow recommendations. Data collected will include anonymised interview transcripts, notes, and conceptual design outputs. The findings will contribute to the development of structured, scalable approaches for managing information complexity in organisational contexts and will be reported in the TU Delft Master's thesis.

ID: 197879

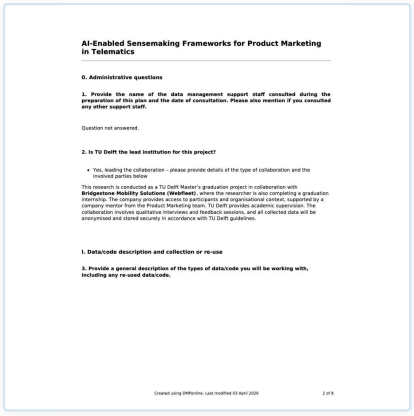
Start date: 11-02-2026

End date: 03-08-2026

Last modified: 03-04-2026

Created using DMPonline. Last modified 03 April 2026 1 of 8

PAGE 1



AI-Enabled Sensemaking Frameworks for Product Marketing in Telematics

0. Administrative questions

1. Provide the name of the data management support staff consulted during the preparation of this plan and the date of consultation. Please also mention if you consulted any other support staff.

Question not answered.

2. Is TU Delft the lead institution for this project?

- Yes, leading the collaboration – please provide details of the type of collaboration and the involved parties below

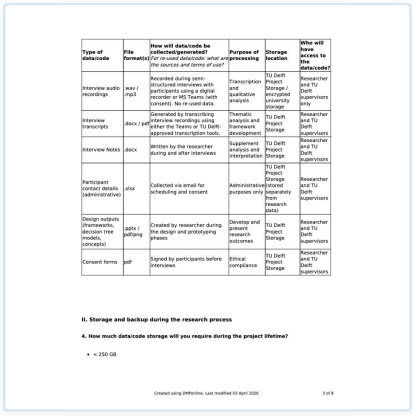
This research is conducted as a TU Delft Master's graduation project in collaboration with **Bridgestone Mobility Solutions (Webfleet)**, where the researcher is also completing a graduation internship. The company provides access to participants and organisational context, supported by a company mentor from the Product Marketing team. TU Delft provides academic supervision. The collaboration involves qualitative interviews and feedback sessions, and all collected data will be anonymised and stored securely in accordance with TU Delft guidelines.

I. Data/code description and collection or re-use

3. Provide a general description of the types of data/code you will be working with, including any re-used data/code.

Created using DMPonline. Last modified 03 April 2026 2 of 8

PAGE 2



Type of data/code	File format(s)	How will data/code be collected/generated? *For re-used data/code: what are the sources and terms of use?*	Purpose of processing	Storage location	Who will have access to the data/code?
Interview audio recordings	.wav / .mp3	Recorded during semi-structured interviews with participants using a digital recorder (with verbal consent). No re-used data.	Transcription and qualitative analysis	TU Delft Project Storage / encrypted university storage	Researcher and TU Delft supervisors only
Interview transcripts	.docx / .pdf	Generated by transcribing interview recordings using either the Teams or TU Delft-approved transcription tools.	Thematic analysis and framework development	TU Delft Project Storage	Researcher and TU Delft supervisors
Interview Notes	.docx	Written by the researcher during and after interviews	Supplement analysis and interpretation	TU Delft Project Storage	Researcher and TU Delft supervisors
Participant contact details (administrative)	.xlsx	Collected via email for scheduling and consent	Administrative purposes only	TU Delft Project Storage stored separately from research data)	Researcher and TU Delft supervisors
Design outputs (frameworks, decision tree models, concepts)	.pptx / .pdf/jpeg	Created by researcher during the design and prototyping phases	Develop and present research outcomes	TU Delft Project Storage	Researcher and TU Delft supervisors
Consent forms	.pdf	Signed by participants before interviews	Ethical compliance	TU Delft Project Storage	Researcher and TU Delft supervisors

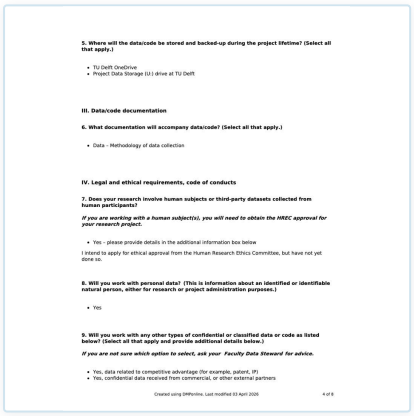
II. Storage and backup during the research process

4. How much data/code storage will you require during the project lifetime?

- < 250 GB

Created using DMPonline. Last modified 03 April 2026 3 of 8

PAGE 3



5. Where will the data/code be stored and backed-up during the project lifetime? (Select all that apply.)

- TU Delft OneDrive
- Project Data Storage (U:) drive at TU Delft

III. Data/code documentation

6. What documentation will accompany data/code? (Select all that apply.)

- Data – Methodology of data collection

IV. Legal and ethical requirements, code of conducts

7. Does your research involve human subjects or third-party datasets collected from human participants?

If you are working with a human subject(s), you will need to obtain the HREC approval for your research project.

- Yes – please provide details in the additional information box below

I intend to apply for ethical approval from the Human Research Ethics Committee, but have not yet done so.

8. Will you work with personal data? (This is information about an identified or identifiable natural person, either for research or project administration purposes.)

- Yes

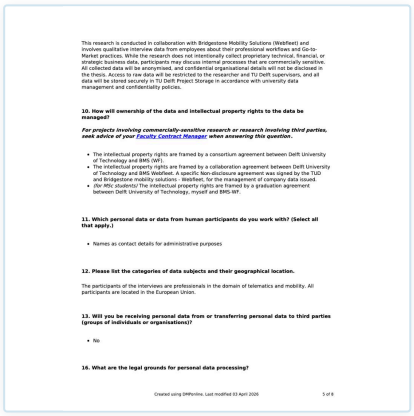
9. Will you work with any other types of confidential or classified data or code as listed below? (Select all that apply and provide additional details below.)

If you are not sure which option to select, ask your Faculty Data Steward for advice.

- Yes, data related to competitive advantage (for example, patent, IP)
- Yes, confidential data received from commercial, or other external partners

Created using DMPonline. Last modified 03 April 2026 4 of 8

PAGE 4



This research is conducted in collaboration with Bridgestone Mobility Solutions (Webfleet) and involves qualitative interview data from employees about their professional workflows and Go-to-Market practices. While the research does not intentionally collect proprietary technical, financial, or strategic business data, participants may discuss internal processes that are commercially sensitive. All collected data will be anonymised, and confidential organisational details will not be disclosed in the thesis. Access to raw data will be restricted to the researcher and TU Delft supervisors, and all data will be stored securely in TU Delft Project Storage in accordance with university data management and confidentiality policies.

10. How will ownership of the data and intellectual property rights to the data be managed?

For projects involving commercially-sensitive research or research involving third parties, seek advice of your Faculty Contract Manager when answering this question.

- The intellectual property rights are framed by a consortium agreement between Delft University of Technology, BMS (WF).
- The intellectual property rights are framed by a collaboration agreement between Delft University of Technology and BMS Webfleet. A specific Non-disclosure agreement was signed by the TUD and Bridgestone mobility solutions - Webfleet, for the management of company data issued.
- *(for MSc students)* The intellectual property rights are framed by a graduation agreement between Delft University of Technology, myself and BMS-WF.

11. Which personal data or data from human participants do you work with? (Select all that apply.)

- Names as contact details for administrative purposes

12. Please list the categories of data subjects and their geographical location.

The participants of the interviews are professionals in the domain of telematics and mobility. All participants are located in the European Union.

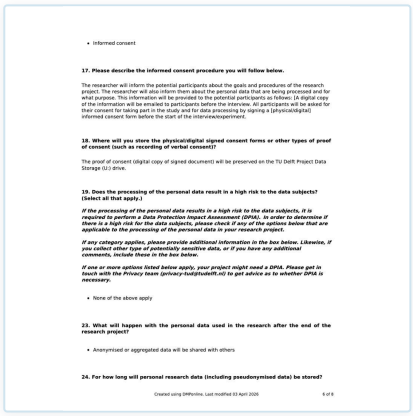
13. Will you be receiving personal data from or transferring personal data to third parties (groups of individuals or organisations)?

- No

16. What are the legal grounds for personal data processing?

Created using DMPonline. Last modified 03 April 2026 5 of 8

PAGE 5



- Informed consent

17. Please describe the informed consent procedure you will follow below.

The researcher will inform the potential participants about the goals and procedures of the research project. The researcher will also inform them about the personal data that are being processed and for what purpose. This information will be provided to the potential participants as follows: [A digital copy of the information will be emailed to participants before the interview. All participants will be asked for their consent for taking part in the study and for data processing by signing a [physical/digital] informed consent form before the start of the interview/experiment.

18. Where will you store the physical/digital signed consent forms or other types of proof of consent (such as recording of verbal consent)?

The proof of consent (digital copy of signed document) will be preserved on the TU Delft Project Data Storage (U:) drive.

19. Does the processing of the personal data result in a high risk to the data subjects? (Select all that apply.)

If the processing of the personal data results in a high risk to the data subjects, it is required to perform a Data Protection Impact Assessment (DPIA). In order to determine if there is a high risk for the data subjects, please check if any of the options below that are applicable to the processing of the personal data in your research project.

If any category applies, please provide additional information in the box below. Likewise, if you collect other type of potentially sensitive data, or if you have any additional comments, include these in the box below.

If one or more options listed below apply, your project might need a DPIA. Please get in touch with the Privacy team (privacy-tud@tudelft.nl) to get advice as to whether DPIA is necessary.

- None of the above apply

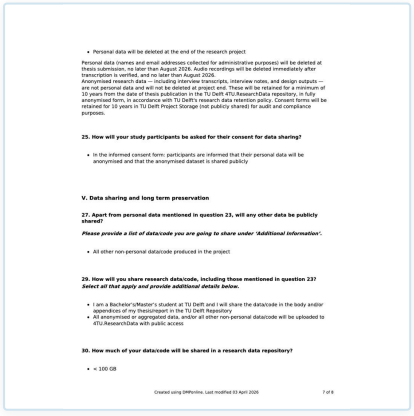
23. What will happen with the personal data used in the research after the end of the research project?

- Anonymised or aggregated data will be shared with others

24. For how long will personal research data (including pseudonymised data) be stored?

Created using DMPonline. Last modified 03 April 2026 6 of 8

PAGE 6



- Personal data will be deleted at the end of the research project

Personal data (names and email addresses collected for administrative purposes) will be deleted at thesis submission, no later than August 2026. Audio recordings will be deleted immediately after transcription is verified, and no later than August 2026.
Anonymised research data — including interview transcripts, interview notes, and design outputs — are not personal data and will not be deleted at project end. These will be retained for a minimum of 10 years from the date of thesis publication in the TU Delft 4TU.ResearchData repository, in fully anonymised form, in accordance with TU Delft's research data retention policy. Consent forms will be retained for 10 years in TU Delft Project Storage (not publicly shared) for audit and compliance purposes.

25. How will your study participants be asked for their consent for data sharing?

- In the informed consent form: participants are informed that their personal data will be anonymised and that the anonymised dataset is shared publicly

V. Data sharing and long term preservation

27. Apart from personal data mentioned in question 23, will any other data be publicly shared?

Please provide a list of data/code you are going to share under 'Additional Information'.

- All other non-personal data/code produced in the project

29. How will you share research data/code, including those mentioned in question 23? ***Select all that apply and provide additional details below.***

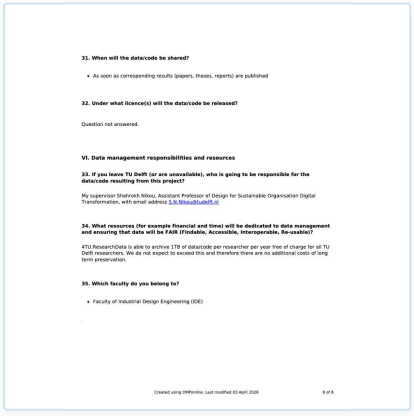
- I am a Bachelor's/Master's student at TU Delft and I will share the data/code in the body and/or appendices of my thesis/report in the TU Delft Repository
- All anonymised or aggregated data, and/or all other non-personal data/code will be uploaded to 4TU.ResearchData with public access

30. How much of your data/code will be shared in a research data repository?

- < 100 GB

Created using DMPonline. Last modified 03 April 2026 7 of 8

PAGE 7



31. When will the data/code be shared?

- As soon as corresponding results (papers, theses, reports) are published

32. Under what licence(s) will the data/code be released?

Question not answered.

VI. Data management responsibilities and resources

33. If you leave TU Delft (or are unavailable), who is going to be responsible for the data/code resulting from this project?

My supervisor Shahrokh Nikou, Assistant Professor of Design for Sustainable Organisation Digital Transformation, with email address S.N.Nikou@tudelft.nl

34. What resources (for example financial and time) will be dedicated to data management and ensuring that data will be FAIR (Findable, Accessible, Interoperable, Re-usable)?

4TU.ResearchData is able to archive 1TB of data/code per researcher per year free of charge for all TU Delft researchers. We do not expect to exceed this and therefore there are no additional costs of long term preservation.

35. Which faculty do you belong to?

- Faculty of Industrial Design Engineering (IDE)

Created using DMPonline. Last modified 03 April 2026 8 of 8

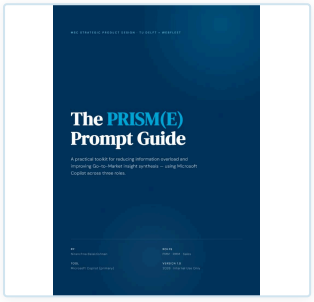
PAGE 8

Appendix C Data management plan — all 8 pages.

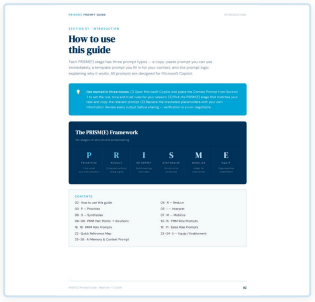
APPENDIX D

The PRISM(E) AI Prompt Guide

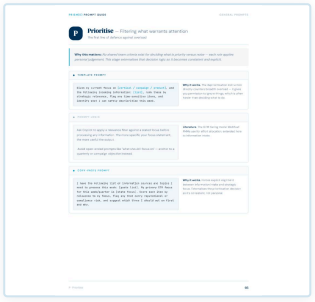
The practitioner-facing prompt guide that operationalises PRISM(E) into ready-to-use AI prompts, co-created and refined with the Product Marketing team, reproduced at reduced size (27 pages).



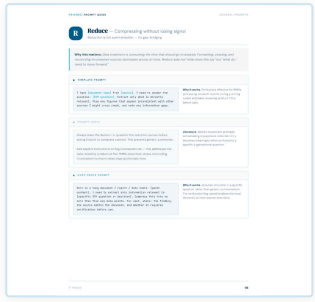
PAGE 1



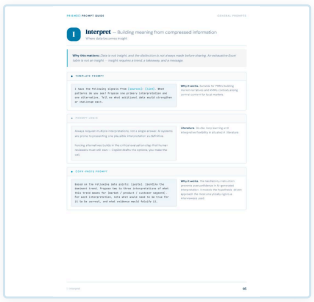
PAGE 2



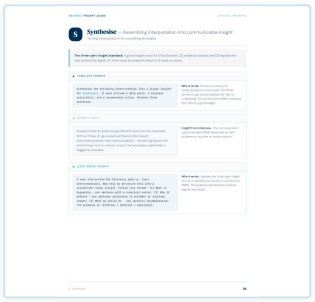
PAGE 3



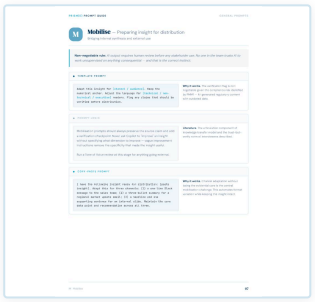
PAGE 4



PAGE 5



PAGE 6



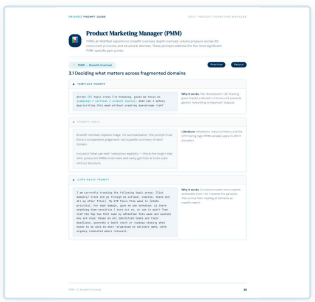
PAGE 7



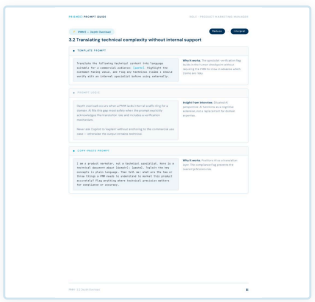
PAGE 8



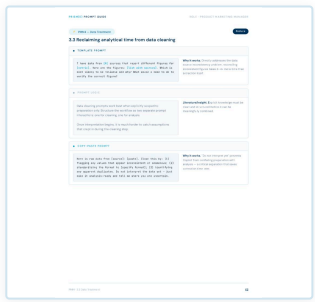
PAGE 9



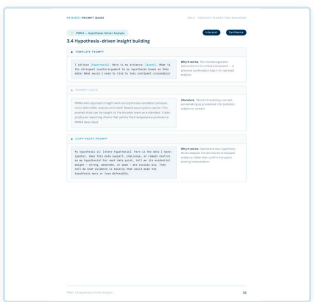
PAGE 10



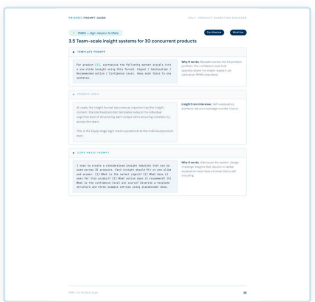
PAGE 11



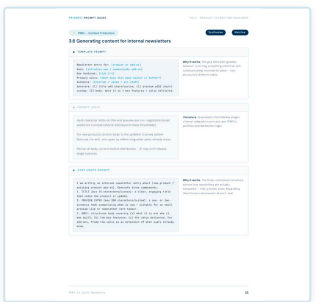
PAGE 12



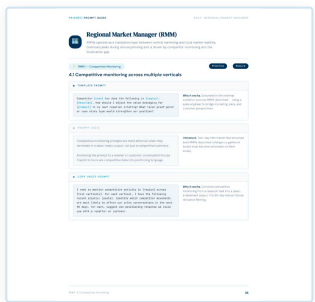
PAGE 13



PAGE 14



PAGE 15



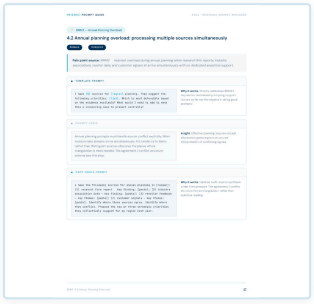
PAGE 16

Appendix D.1 The PRISM(E) AI Prompt Guide — pages 1–16 of 27.

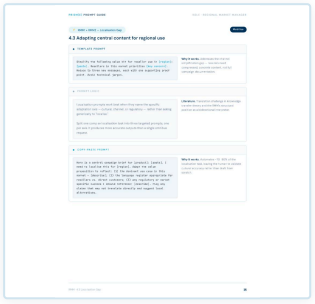
APPENDIX D · CONTINUED

The PRISM(E) AI Prompt Guide

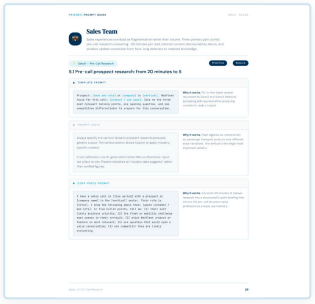
Continued — pages 17–27 of 27.



PAGE 17



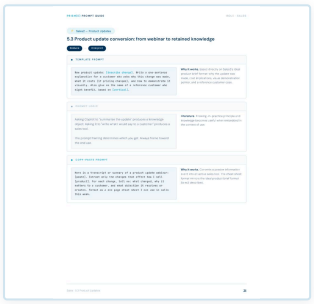
PAGE 18



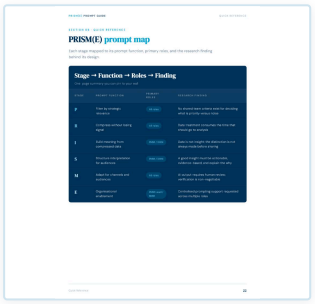
PAGE 19



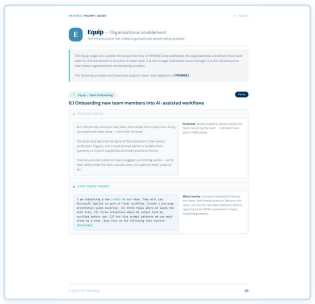
PAGE 20



PAGE 21



PAGE 22



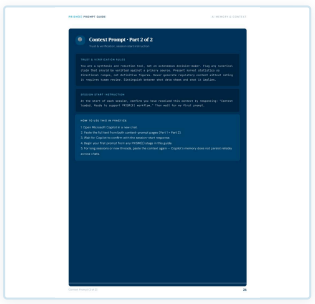
PAGE 23



PAGE 24



PAGE 25



PAGE 26



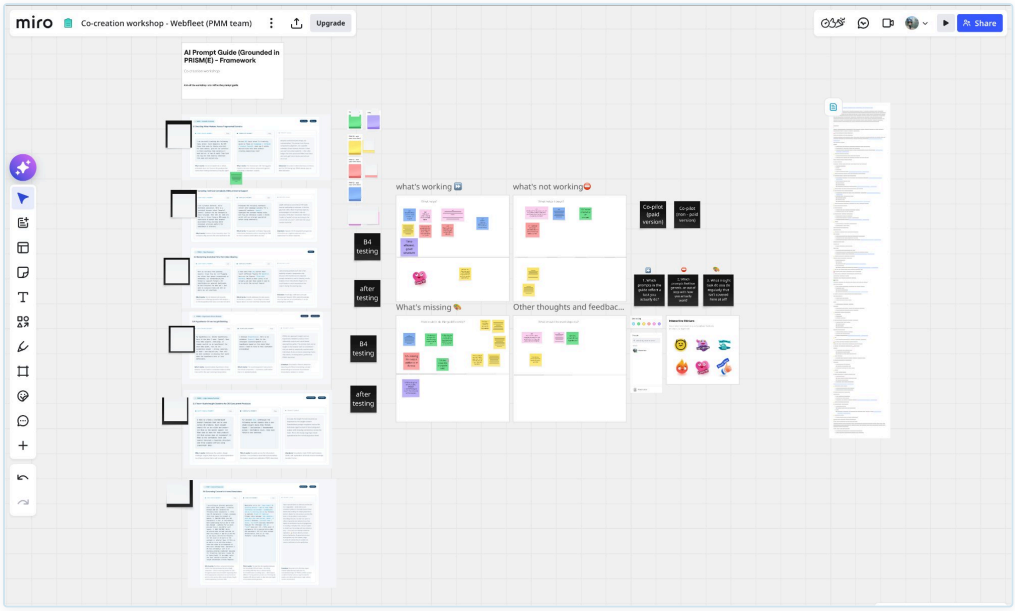
PAGE 27

Appendix D.2 The PRISM(E) AI Prompt Guide — pages 17–27 of 27.

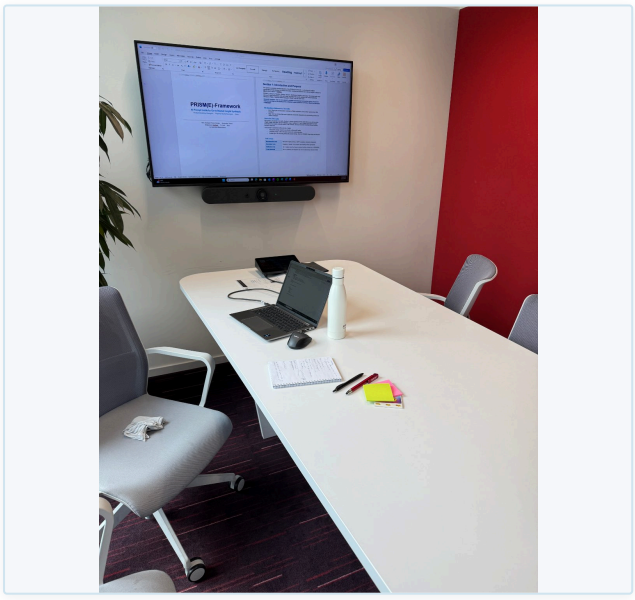
APPENDIX E

Co-creation workshop materials

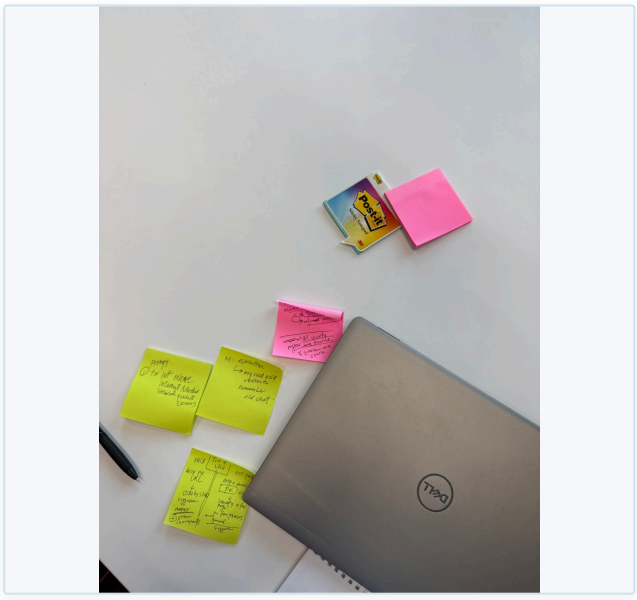
Working materials from the co-creation sessions that refined the PRISM(E) AI Prompt Guide with practitioners — the shared Miro board and photographs from the pilot session.



MIRO BOARD · CO-CREATION WORKSHOP WITH THE PRODUCT MARKETING TEAM



PILOT SESSION · GUIDE REVIEWED ON SCREEN



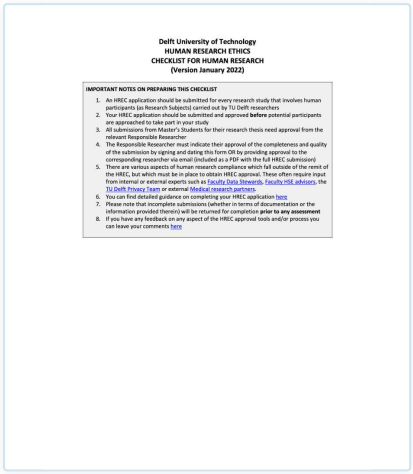
PILOT SESSION · PRACTITIONER ANNOTATIONS

Appendix E Co-creation workshop board (Miro) used to refine the prompt guide with the PMM team — what's working, what's not, what's missing, and next steps — alongside photographs from the pilot co-creation session.

APPENDIX F

HREC

The Human Research Ethics Committee (HREC) submission, approval, and accompanying documents for the study, reproduced at reduced size (pages 1–6 of 11).

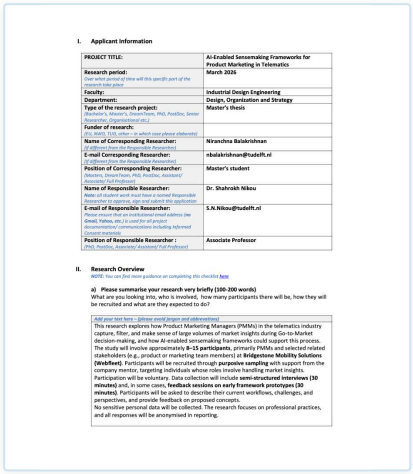


Delft University of Technology
HUMAN RESEARCH ETHICS
CHECKLIST FOR HUMAN RESEARCH
(Version January 2022)

IMPORTANT NOTES ON PREPARING THIS CHECKLIST

1. An HREC application should be submitted for every research study that involves human participants (as Research Subjects) carried out by TU Delft researchers
2. Your HREC application should be submitted and approved **before** potential participants are approached to take part in your study
3. All submissions from Master's Students for their research thesis need approval from the relevant Responsible Researcher
4. The Responsible Researcher must indicate their approval of the completeness and quality of the submission by signing and dating this form OR by providing approval to the corresponding researcher via email (included as a PDF with the full HREC submission)
5. There are various aspects of human research compliance which fall outside of the remit of the HREC, but which must be in place to obtain HREC approval. These often require input from internal or external experts such as Faculty Data Stewards, Faculty HSE advisors, the TU Delft Privacy Team or external Medical research partners.
6. You can find detailed guidance on completing your HREC application here
7. Please note that incomplete submissions (whether in terms of documentation or the information provided therein) will be returned for completion **prior to any assessment**
8. If you have any feedback on any aspect of the HREC approval tools and/or process you can leave your comments here

PAGE 1



I. Applicant Information

PROJECT TITLE:	**AI-Enabled Sensemaking Frameworks for Product Marketing in Telematics**
Research period: *Over what period of time will this specific part of the research take place*	**March 2026**
Faculty:	**Industrial Design Engineering**
Department:	**Design, Organization and Strategy**
Type of the research project: *(Bachelor's, Master's, Graduate Thesis, PhD, PostDoc, Senior Researcher, Organizational etc.)*	**Master's thesis**
Funder of research: *(e.g. NWO, TUD, other) – in which case please elaborate)*	
Name of Corresponding Researcher: *(if different from the Responsible Researcher)*	**Niranchna Balakrishnan**
E-mail Corresponding Researcher: *(if different from the Responsible Researcher)*	**nbalakrishnan@tudelft.nl**
Position of Corresponding Researcher: *(Master's, Graduate Thesis, PhD, PostDoc, Assistant/ Full Professor etc.)*	**Master's student**
Name of Responsible Researcher: *Note: all student work must have a named Responsible Researcher to approve, sign and submit this application*	**Dr. Shahrokh Nikou**
E-mail of Responsible Researcher: *Please ensure that an institutional email address (no Gmail, Yahoo, etc.) is used for all project documentation/ communications including Informed Consent materials*	**S.N.Nikou@tudelft.nl**
Position of Responsible Researcher : *(PhD, PostDoc, Associate/ Assistant/ Full Professor)*	**Associate Professor**

II. Research Overview

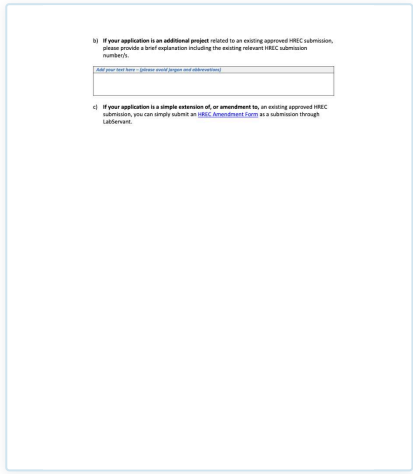
NOTE: You can find more guidance on completing this checklist here

a) Please summarise your research very briefly (100-200 words)
What are you looking into, who is involved, how many participants there will be, how they will be recruited and what are they expected to do?

Add your text here – (please avoid jargon and abbreviations)

This research explores how Product Marketing Managers (PMMs) in the telematics industry capture, filter, and make sense of large volumes of market insights during Go-to-Market decision-making, and how AI-enabled sensemaking frameworks could support this process. The study will involve approximately **8–15 participants**, primarily PMMs and selected related stakeholders (e.g., product or marketing team members) at **Bridgestone Mobility Solutions (WebText)**. Participants will be recruited through **purposive sampling** with support from the company mentor, targeting individuals whose roles involve handling market insights. Participation will be voluntary. Data collection will include **semi-structured interviews (30 minutes)** and, in some cases, **feedback sessions on early framework prototypes (30 minutes)**. Participants will be asked to describe their current workflows, challenges, and perspectives, and provide feedback on proposed concepts.
No sensitive personal data will be collected. The research focuses on professional practices, and all responses will be anonymised in reporting.

PAGE 2

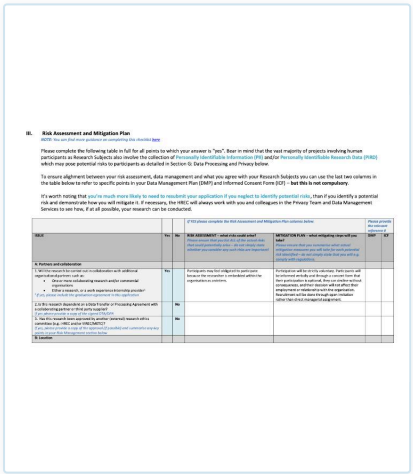


b) **If your application is an additional project** related to an existing approved HREC submission, please provide a brief explanation including the existing relevant HREC submission number/s.

Add your text here – (please avoid jargon and abbreviations)

c) **If your application is a simple extension of, or amendment to,** an existing approved HREC submission, you can simply submit an HREC Amendment Form as a submission through LabServant.

PAGE 3



III. Risk Assessment and Mitigation Plan

NOTE: You can find more guidance on completing this checklist here

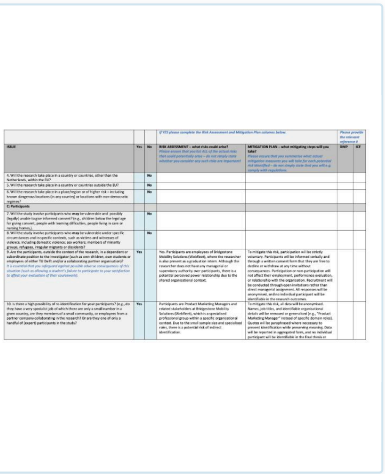
Please complete the following table in full for all points to which your answer is "yes". Bear in mind that the vast majority of projects involving human participants as Research Subjects also involve the collection of Personally Identifiable Information (PII) and/or Personally Identifiable Research Data (PIRD) which may pose potential risks to participants as detailed in Section G: Data Processing and Privacy below.

To ensure alignment between your risk assessment, data management and what you agree with your Research Subjects you can use the last two columns in the table below to refer to specific points in your Data Management Plan (DMP) and Informed Consent Form (ICF) – **but this is not compulsory.**

It's worth noting that you're much more likely to need to resubmit your application if you neglect to identify potential risks, than if you identify a potential risk and demonstrate how you will mitigate it. If necessary, the HREC will always work with you and colleagues in the Privacy Team and Data Management Services to see how, if at all possible, your research can be conducted.

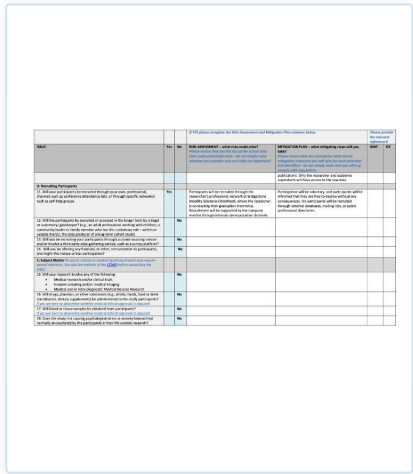
			If YES please complete the Risk Assessment and Mitigation Plan columns below		Please provide the relevant reference #
ISSUE	**Yes**	**No**	**RISK ASSESSMENT – what risks could arise?** *Please ensure that you systematise what certain mitigation measures you will take for each potential risk. Please note that even if you have already stated that you consider any such risks are important*	**MITIGATION PLAN – what mitigating steps will you take?** *Please ensure that you systematise what certain mitigation measures you will take for each potential risk. Please note that even if you have already stated that you consider any such risks are important*	**DMP**	**ICF**
A. Partners and collaboration						
1. Will the research take place in a country or countries, other than the Netherlands, outside the EU?		No				
2. Will the research take place in a country or countries outside the EU?		No				
3. Will the research take place in a place/region of higher risk in terms of known dangerous locations (or any country or location with non-democratic rule)?		No				
B. Participants						
1. Will the study involve participants who may be cognitively and socially (legally) unable to give informed consent? (e.g., children below the legal age for giving consent, people with learning difficulties, people living in care or nursing homes)		No				
2. Will the study involve participants who may be vulnerable and/or care for, or live with, vulnerable people? (e.g., children, people with learning disabilities, elderly, people living in care or nursing homes, refugees, irregular migrants, or prisoners)		No				
3. Are the participants outside the scope of the research, in a dependency or unequal relationship with the researcher? (e.g., students and staff, own students or employees of either TU Delft and/or a collaborating partner organisation)	Yes		Yes. Participants are employees of Bridgestone Mobility Solutions (WebText), whom the researcher is also present as a guide and an intern. Although the researcher does not have any managerial or supervisory role over participants, there is a potential perceived power relationship due to the shared organisational context.	To mitigate this risk, participation will be strictly voluntary. Participants will be informed verbally and through a written consent form that they are free to decline or withdraw at any time without consequence. The researcher will explicitly state that not effect their employment, performance evaluations, or relationship with the organisation. Recruitment will be conducted through email invitations rather than in-person requests, and no one from the participant's management, and no individual participant will be informed about who has been invited.		
4. Is there a high possibility of re-identification for your participants? (e.g., do they have a very specific job of which there are only a small number in a given country, are they members of a small community, or employees from a small organisation, and there is a risk that a combination of the data of only a handful of (event) participants in the study?)		Yes		To mitigate this risk, all data will be anonymised. Names, job titles, and identifiable organisational details will be removed or generalised (e.g., "Product Marketing Manager, Europe" instead of specific company and location). Quotes will be paraphrased where necessary to prevent identification while preserving meaning. Data will be stored securely, and only the researcher and supervisors will be identifiable to the three thesis		

PAGE 4



			If YES please complete the Risk Assessment and Mitigation Plan columns below		Please provide the relevant reference #
ISSUE	**Yes**	**No**	**RISK ASSESSMENT – what risks could arise?** *Please ensure that you systematise what certain mitigation measures you will take for each potential risk. Please note that even if you have already stated that you consider any such risks are important*	**MITIGATION PLAN – what mitigating steps will you take?** *Please ensure that you systematise what certain mitigation measures you will take for each potential risk. Please note that even if you have already stated that you consider any such risks are important*	**DMP**	**ICF**
5. Will the research take place in a country or countries, other than the Netherlands, outside the EU?		No				
6. Will the research take place in a country or countries outside the EU?		No				
7. Will the research take place in a place/region of higher risk in terms of known dangerous locations (or any country or location with non-democratic rule)?		No				
C. Participants						
1. Will the study involve participants who may be cognitively and socially (legally) unable to give informed consent? (e.g., children below the legal age for giving consent, people with learning difficulties, people living in care or nursing homes, refugees, irregular migrants, or prisoners)		No				
2. Will the study involve participants who may be vulnerable and/or care for, or live with, vulnerable people? (e.g., children, people with learning disabilities, elderly, people living in care or nursing homes, refugees, irregular migrants or prisoners)		No				
3. Are the participants outside the scope of the research, in a dependency or unequal relationship with the researcher? (e.g., students and staff, own students or employees of either TU Delft and/or a collaborating partner organisation)	Yes		Yes. Participants are employees of Bridgestone Mobility Solutions (WebText), whom the researcher is also present as a guide and an intern. Although the researcher does not have any managerial or supervisory role over participants, there is a potential perceived power relationship due to the shared organisational context.	To mitigate this risk, participation will be strictly voluntary. Participants will be informed verbally and through a written consent form that they are free to decline or withdraw at any time without consequence. The researcher will explicitly state that not effect their employment, performance evaluations, or relationship with the organisation. Recruitment will be conducted through email invitations rather than in-person requests, and no one from the participant's management, and no individual participant will be informed about who has been invited.		
4. Is there a high possibility of re-identification for your participants? (e.g., do they have a very specific job of which there are only a small number in a given country, are they members of a small community, or employees from a small organisation, and there is a risk that a combination of the data of only a handful of (event) participants in the study?)		Yes		To mitigate this risk, all data will be anonymised. Names, job titles, and identifiable organisational details will be removed or generalised (e.g., "Product Marketing Manager, Europe" instead of specific company and location). Quotes will be paraphrased where necessary to prevent identification while preserving meaning. Data will be stored securely, and only the researcher and supervisors will be identifiable to the three thesis		

PAGE 5



			If YES please complete the Risk Assessment and Mitigation Plan columns below		Please provide the relevant reference #
ISSUE	**Yes**	**No**	**RISK ASSESSMENT – what risks could arise?** *Please ensure that you systematise what certain mitigation measures you will take for each potential risk. Please note that even if you have already stated that you consider any such risks are important*	**MITIGATION PLAN – what mitigating steps will you take?** *Please ensure that you systematise what certain mitigation measures you will take for each potential risk. Please note that even if you have already stated that you consider any such risks are important*	**DMP**	**ICF**
				publications. Only the researcher and academic supervisors will have access to the raw data.		
D. Recruiting Participants						
11. Will your participants be recruited through your own, professional, channels such as conference attendance lists, or through specific networks such as self-help groups	Yes		Participants will be recruited through the researcher's professional network at Bridgestone Mobility Solutions (WebText), where the researcher is also present as a guide and an intern. Recruitment will be supported by the company mentor leveraging internal communication channels.	Invited via email, who voluntary, and will contain: will be informed that they are free to decline without any consequences. No participants will be recruited during working hours, meeting time, or public professional time slots.		
12. Will the participants be recruited or accessed in the longer term by a target or customary gatekeeper? (e.g., an adult performance working with children in community leader or family members who live this customary role – within Western society, this is usually parents for their children)		No				
13. Will you be recruiting your participants through an email recruiting service and/or involve a third party data gathering service, such as a survey platform?		No				
14. Will you be using any incentives, remuneration or compensation to participants, and might this involve or link participation?		No				
E. Subject Matter *Standard statistics on medical (particularly health) may require ethical clearance. Be aware the website of the CCMO makes interesting the topic!*						
15. Will your research involve any of the following:		No				
• Drugs, pharmaceuticals or medical devices • Invasive sampling and/or medical imaging • Medical aid in Anthropometric Medical Devices Research						
16. Will drugs, placebos, or other substances (e.g., drinks, foods, food or dietary supplements, dietary regimes) be administered to the study participants?		No				
17. Will blood or tissue samples be obtained from participants?		No				
18. Will the research involve the analysis of blood, genetic material, tissue or other bodily samples? Will you be working with (anonymised) human material or data originally encountered by the participants in their life outside research?		No				

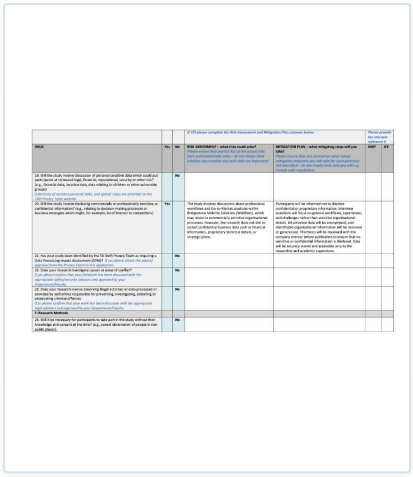
PAGE 6

Appendix F.1 HREC approval and revisions — pages 1–6 of 11.

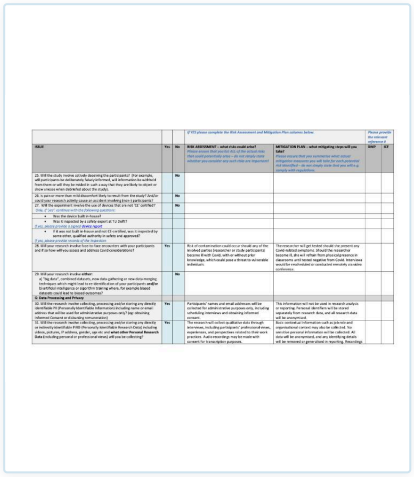
APPENDIX F · CONTINUED

HREC

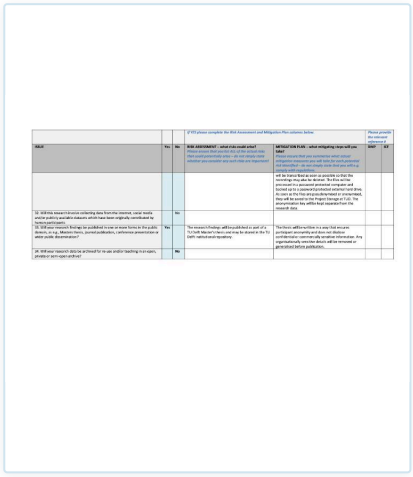
Continued — pages 7–11 of 11.



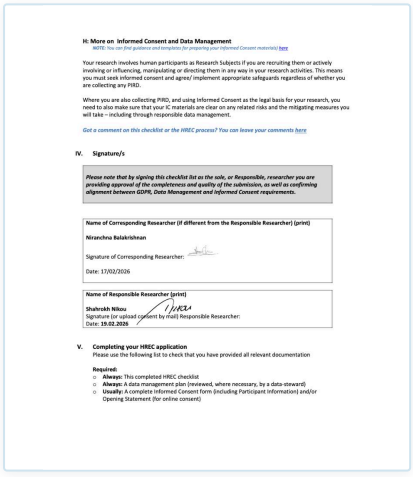
PAGE 7



PAGE 8



PAGE 9



H: More on Informed Consent and Data Management

Your research involves human participants as Research Subjects if you are recruiting them or actively involving or influencing, manipulating or directing them in any way in your research activities. This means you must seek informed consent and agree/ implement appropriate safeguards regardless of whether you are collecting any PIRD.

Where you are also collecting PIRD, and using Informed Consent as the legal basis for your research, you need to also make sure that your IC materials are clear on any related risks and the mitigating measures you will take – including through responsible data management.

Got a comment on this checklist or the HREC process? You can leave your comments here

IV. Signature/s

Please note that by signing this checklist list as the sole, or Responsible, researcher you are providing approval of the completeness and quality of the submission, as well as confirming alignment between GDPR, Data Management and Informed Consent requirements.

Name of Corresponding Researcher (if different from the Responsible Researcher) (print)

Niranchna Balakrishnan

Signature of Corresponding Researcher:

Date: 17/02/2026

Name of Responsible Researcher (print)

Shahrokh Nikou

Signature (or upload consent by mail) Responsible Researcher:

Date: **19.02.2026**

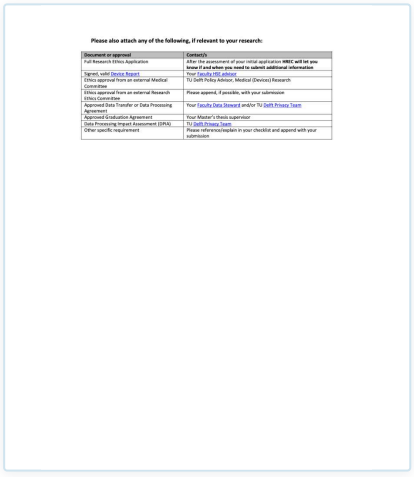
V. Completing your HREC application

Please use the following list to check that you have provided all relevant documentation

Required:

- **Always:** This completed HREC checklist
- **Always:** A data management plan (reviewed, where necessary, by a data-steward)
- **Usually:** A complete Informed Consent form (including Participant Information) and/or Opening Statement (for online consent)

PAGE 10



Please also attach any of the following, if relevant to your research:

Document or approval	Contact/s
Full Research Ethics Application	After the assessment of your initial application **HREC will let you know if and when you need to submit additional information**
Signed, valid Device Report	Your Faculty HSE advisor
Ethics approval from an external Medical Committee	TU Delft Policy Advisor, Medical (Devices) Research
Ethics approval from an external Research Ethics Committee	Please append, if possible, with your submission
Approved Data Transfer or Data Processing Agreement	Your Faculty Data Steward and/or TU Delft Privacy Team
Approved Graduation Agreement	Your Master's thesis supervisor
Data Processing Impact Assessment (DPIA)	TU Delft Privacy Team
Other specific requirement	Please reference/explain in your checklist and append with your submission

PAGE 11

Appendix F.2 HREC approval and revisions – pages 7–11 of 11.

THESIS · STRATEGIC ROADMAP · THREE-HORIZON PLAN

From pilot to permanent capability — *embedding PRISM(E) into Product Marketing*

v 1.0 · 18-month plan
Owner · PMM Lead
Primary tool · Microsoft Copilot
Date · May 2026

HORIZON 1 FOUNDATION

Foundation & Pilot

Months 01 – 04

Establishing infrastructure and validating PRISM(E) in controlled conditions with a small PMM cohort, producing the first AI-assisted GTM insight outputs.

HORIZON 2 SCALE

Scale & Integration

Months 05 – 10

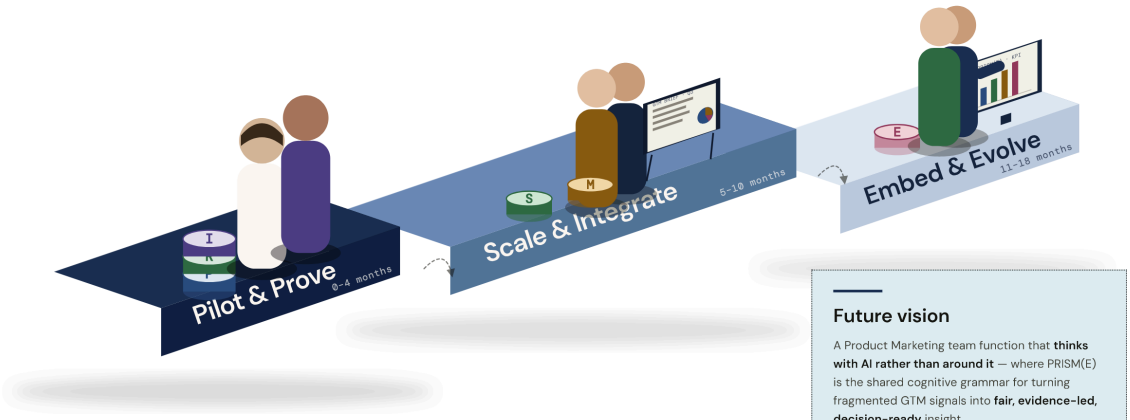
Rolling PRISM(E) across PMM and Regional Market Managers, integrating with the quarterly GTM cycle and formalising the Synthesise-to-Mobilise handoff.

HORIZON 3 EMBED

Institutionalise & Evolve

Months 11 – 18

Embedding PRISM(E) as a permanent organisational capability, measuring long-term impact on GTM quality and iterating the framework on lived practice.



Future vision

A Product Marketing team function that **thinks with AI rather than around it** — where PRISM(E) is the shared cognitive grammar for turning fragmented GTM signals into **fair, evidence-led, decision-ready** insight.

KEY MILESTONES

M02 ··· M18 · eighteen-month timeline

M02 Pilot cohort onboarded · baseline measured

M04 First AI-assisted briefs · H1 review

M06 Full PMM + RMM rollout complete

M09 Embedded in quarterly GTM cycle

M12 Longitudinal evaluation begins · v2 drafted

M18 Institutionalised · expansion scoped

E

Equip — *runs across all three horizons*

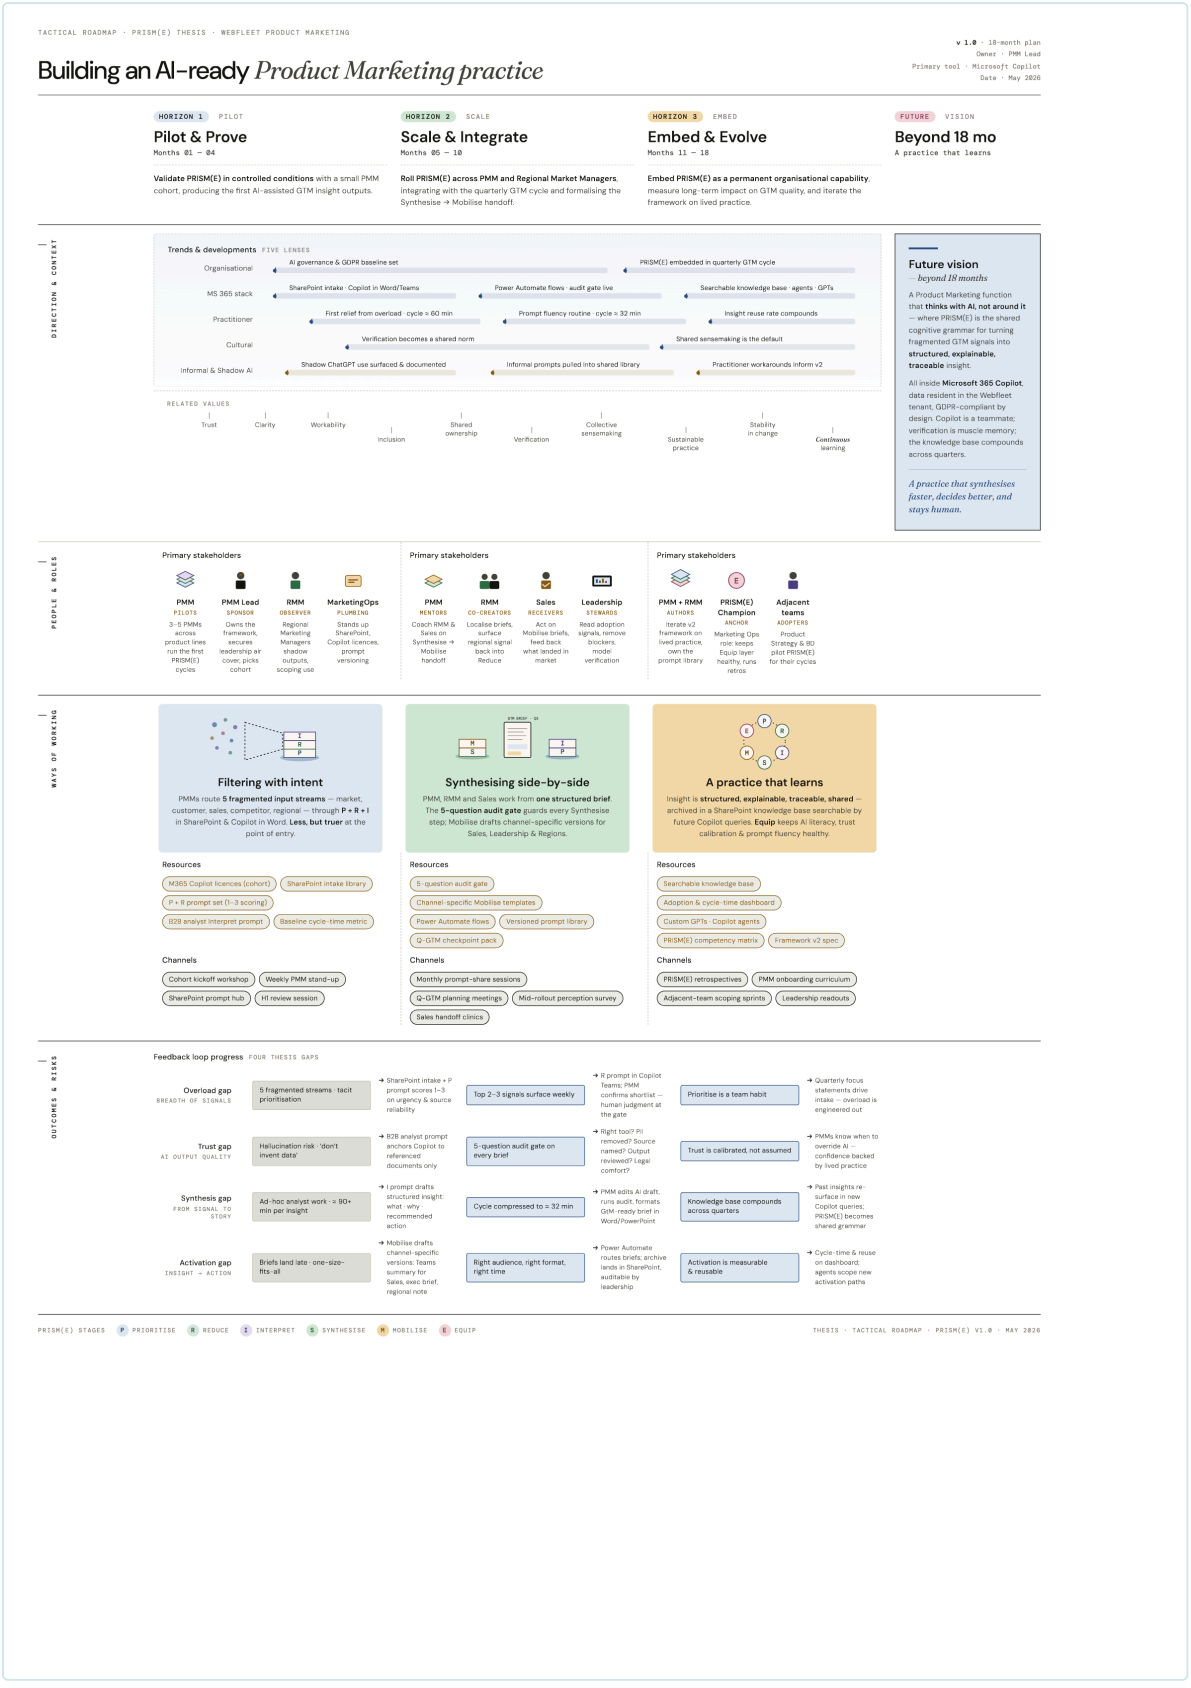
The sixth PRISM(E) element is not sequential. Organisational enablement must be sustained continuously: AI governance, GDPR compliance, leadership sponsorship and capability-building are prerequisites *and* ongoing conditions, not one-time activities. Without a healthy Equip layer, all five sequential stages degrade.

- AI governance policy
- GDPR & data handling
- Leadership sponsorship
- Copilot licence management
- Prompt literacy training
- Change communication

PRISM(E) STAGES — P PRIORITISE · R REDUCE · I INTERPRET · S SYNTHESISE · M MOBILISE · E EQUIP

THESIS · STRATEGIC ROADMAP · WEBFLEET PRODUCT MARKETING — PRISM(E) · V1.0 · MAY 2026

Appendix G.1 Strategic roadmap, full-size reference (see Fig. 8.3, §8.2) — from pilot to permanent capability across three horizons.



Appendix G.2 Tactical roadmap, full-size reference (see Fig. 8.4, §8.3) — building an AI-ready practice through five lenses.

PRISM(E)

"AI helps most not when it makes the decision, but when it gives people better material to decide with".

— NIRANCHNA BALAKRISHNAN

MSc Strategic Product Design
Faculty of Industrial Design Engineering
Delft University of Technology

IN COLLABORATION WITH WEBFLEET - A BRIDGESTONE COMPANY
AMSTERDAM ZUID ·JUNE 2026