August 2025

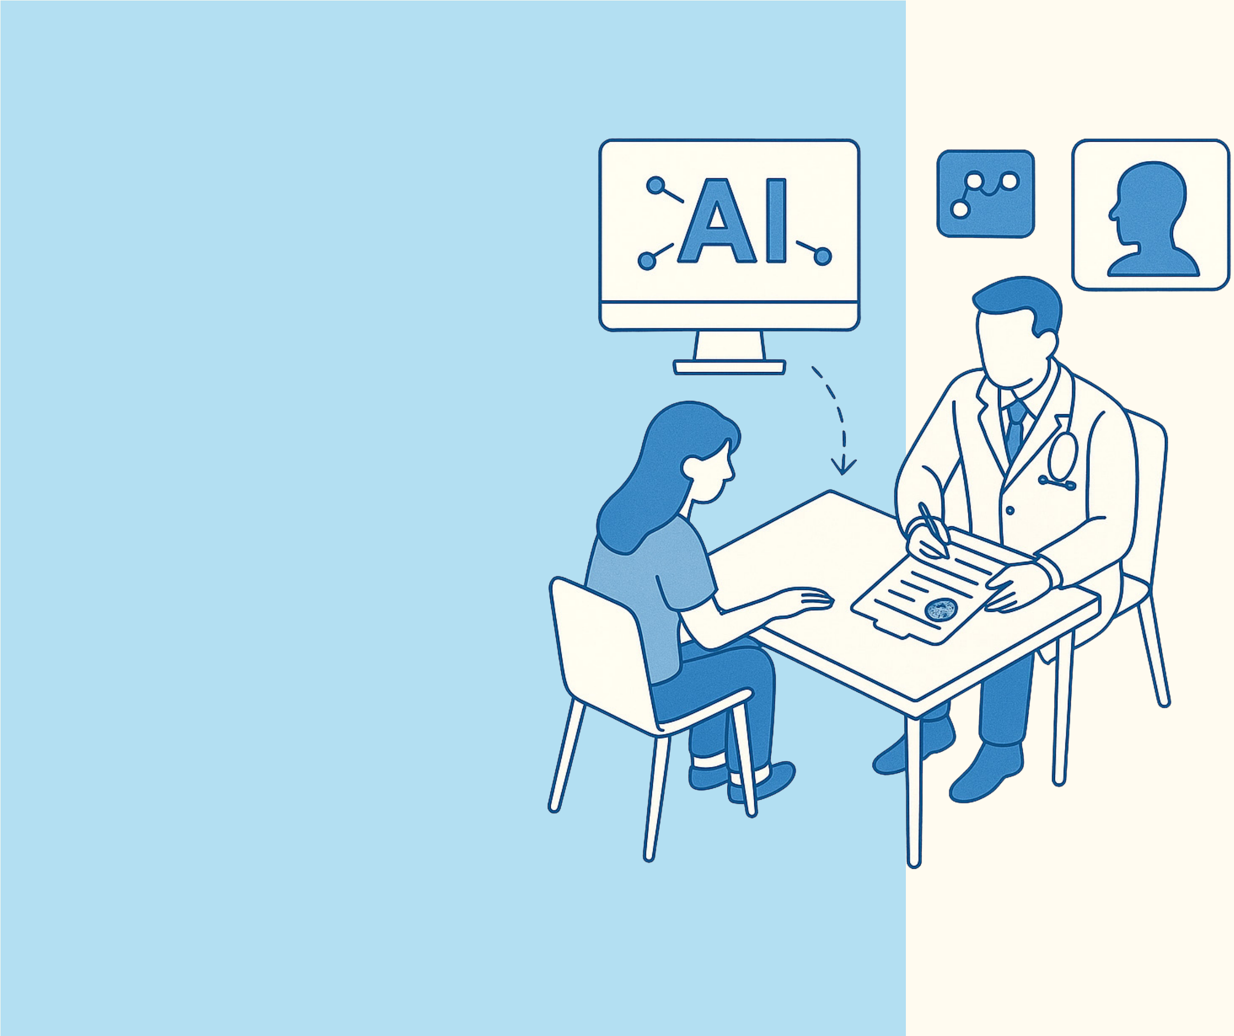

Reimagining Carepathways

An AI driven approach to personalizing chronic care at scale

Gautham Ravikiran

MSc. Design for interaction

Master Thesis

Reimagining Carepathways

An AI driven approach to personalizing chronic care at scale

Gautham Ravikiran
5921171
MSc Design for Interaction - Medisign specialisation

Defended on the *19th of August 2025*

Delft University of technology
Faculty of Industrial design engineering
Landbergstraat 15
2628 CE Delft

Supervisory team

Chair
Prof. dr. ir. R.H.M. Goossens,
Human Centered Design (TU Delft)

Mentors
Dr. ir. J. Jung
Design, Organisation and Strategy (TU Delft)

Dr. ir. H. Torkamaan
Systems Engineering, Multi Actor Systems (TU Delft)

A project in collaboration with

Flagship Consultation Room of the future
Convergence Health & Technology
Erasmus Medical Center,
Rotterdam

Executive Summary

This thesis investigates how unstructured patient narratives can inform the design of digital health tools that incorporate behavioral insights into chronic care management. The project focuses on three specific conditions: colorectal cancer, pulmonary fibrosis, and sarcoidosis. It aims to bridge the gap between patient-reported experiences and clinical decision-making by developing data-driven systems that enhance personalization and support across both patient and clinician workflows.

The research process began with the collection and preprocessing of over 13,000 anonymized forum posts from disease-specific online communities. These texts were analyzed using GPT-4o-assisted semantic clustering to extract latent behavioral themes. We further built a multi-stage verification pipeline combining chain-of-verification prompting, behavioral science frameworks, and manual review to generate rich patient behavior profiles.

Each patient profile captured a specific pattern of goals, motivations, challenges and narratives that were subsequently validated through expert review sessions with clinicians from Erasmus Medical Center. The behaviors were further mapped onto disease-specific care pathways derived from real clinical data and treatment guidelines.

Informed by these behavioral profiles, the design process produced two complementary interventions. The first is a patient-facing mobile application that uses clinically recognized tools such as PAM (Patient Activation Measure) to support patient goal setting focused on quality of life. The app enables mood logging, symptom journaling and facilitates symbolic peer support through real-time emotion-matching.

The second intervention is a clinician-facing decision-support interface that can be embedded within the HIX electronic health record system used across hospitals in the Netherlands. This dashboard synthesizes clinical and behavioral data into a unified view and provides AI-generated care modules that can be reviewed, modified, and assigned by clinicians within existing workflows.

The project resulted in a reproducible data-to-design pipeline for generating patient profiles from unstructured narrative data. It delivered eight validated cross-condition patient archetypes, high-fidelity prototypes for both patient and clinician interfaces, and a Vision 2040 roadmap for the future of personalized care planning. The roadmap outlines a phased implementation strategy with clear milestones for pilot deployment, system integration, and national adoption, grounded in existing healthcare policy and infrastructure constraints.

This thesis demonstrates that patient-authored narratives contain clinically relevant behavioral information. By transforming these narratives into structured archetypes and embedding them into decision-support tools, the project provides a scalable framework for advancing personalized, behavior-aware chronic care delivery.

Acknowledgements

To my chair, **Richard**, thank you for taking me on and creating a space where I could explore ideas freely. Your steady support and thoughtful guidance helped me stay grounded when the project became overwhelming. Your calm words of encouragement gave me confidence during moments of uncertainty, and I always left our conversations with greater clarity and focus.

To **Jiwon**, my mentor, thank you for being a constant source of inspiration and challenge. Your enthusiasm from day one gave me the drive to push forward and aim higher. You pushed me to think critically, build with intention, and back up every idea with evidence. Your sharp questions, honest feedback, and unwavering support helped me grow as a designer and researcher. You believed in the ambition of this work and helped me see how far it could go.

To **Helma**, thank you for guiding me through the technical parts of this project. You pushed me to speak more precisely, structure my methods clearly, and strengthen the technical backbone of my work. Your attention to detail and insistence on clarity helped me sharpen both my thinking and communication.

To the **clinicians** who took part in this project, thank you for your time, insights, and curiosity. Your willingness to engage with this work and reflect critically on its value motivated me to do my best. This project would not have been possible without your input and openness to new possibilities in care.

To the **HCD-IO and Erasmus MC secretaries**, thank you for your help in coordinating key meetings, booking rooms, and making sure everything ran smoothly behind the scenes. Your support made many important conversations possible.

To my **parents**, thank you for your unwavering emotional support. Your strength held me up during moments when my mental strength was being tested, and I am grateful to have had you by my side, even from afar.

To my **friends**, thank you for your support, creativity, and constant encouragement. You challenged me to think in new directions and reminded me to keep going when things got tough. Your presence helped keep the process joyful and ambitious.

This thesis is the result of collective effort, trust, and generosity from many people. I feel incredibly lucky to have had each of you on this journey.

Guide for the reader

This thesis presents an interdisciplinary project that combines LLM analysis, clinical insight, and behavior-centered design to improve chronic care. It follows a structured flow from the collection and processing of unstructured patient data to the generation of behavioral archetypes, the development of digital health interface, and validation through expert clinical feedback.

Chapter 1 introduces the motivation, central research questions, and scope. Chapters 3 and 4 explain the computational process, including the use of NLP and large language models to identify behavioral patterns and generate patient archetypes. Chapters 5 and 6 translate these findings into clinically relevant insights. Chapters 7 and 8 present the Vision 2040 roadmap and two design outcomes: a patient-facing mobile app and a clinician-facing dashboard. These chapters are especially useful for readers interested in service innovation, interface design, and health system integration. Chapter 9 documents feedback from healthcare professionals and highlights the conditions for real-world implementation. Chapters 2 and 10 provide theoretical framing and final reflections.

The appendix includes analyzed data, results from computational analysis and other workflow diagrams. Readers can move through the thesis sequentially or explore sections based on interest.

Key abbreviations used throughout the thesis:

AI – Artificial Intelligence
ML – Machine Learning
LLM – Large Language Model
NLP – Natural Language Processing
PAM – Patient Activation Measure
PROM – Patient-Reported Outcomes Measures
QoL – Quality of Life
EHR – Electronic Health Record
HCP – Healthcare Professional
HIX – Electronic Health Record platform
Co-Ve – Chain-of-Verification

Contents

Chapter 1 Background

Technologies such as artificial intelligence, machine learning, and natural language processing are increasingly reshaping healthcare practices and the way patient care is conceptualized. Leveraging large-scale, patient-generated content such as blogs, forums, and social platforms, present an opportunity to uncover deeper insights into patient experiences.

This project is based on Dr. Ir. Jiwon Jung's research: "Developing Data-enabled Design in the field of digital Health" (Jung, 2023). Within this research area, this master thesis builds on the previous efforts of Master's Graduation students who collaborated with Dr. Ir. Jiwon Jung to visualize patient community journey maps and ideate on strategies and design interventions for three patient communities.

This research focuses on three online patient communities: colorectal cancer, pulmonary fibrosis, and sarcoidosis. These communities are hosted on peer support platforms established by healthcare institutions and patient advocacy groups. They bring together individuals closely involved in the care journey of these conditions involving patients themselves and family members, partners, and occasionally healthcare professionals. Within these spaces, members exchange information, ask questions, reflect on treatments, and share personal experiences. Previous masters' thesis has analysed these online patient communities individually and all together to identify commonalities and develop appropriate strategies and solutions to cater to these needs (Kicken, 2023; Sun, 2023; Zhang, 2023; Quijada Fernández, 2024).

The current master thesis aims to further analyse first-hand patient data to design dynamic and intelligent digital care pathways to help healthcare providers create personalized care plans for patients suffering from chronic conditions. The project will also explore a new method of evidence-based patient behavior profiling to create scalable support strategies that can also be extended into other chronic conditions.

1.1 The problem context

The healthcare system is under pressure. The world's population is ageing while at the same time, there is shortage of healthcare staff for care delivery. With the rising demand for care and approaching crisis of personnel, there is a rapidly increasing accessibility gap that is being created. It is important now, more than ever, to address this societal challenge (Gupta Strategists,2021; Winter et al, 2020).

Previous studies show that chronic patients incur multifaceted needs throughout their care journey. Apart from clinical care needs, the complex care needs of chronic patients are dependent on five dimensions of health that includes physical, psychological, social, intellectual and spiritual well-being (J. R. Hjelm, 2009; Iglesias et al., 2018). These need dimensions are further dependent on multiple factors and occur at varying points in the patient journey making it difficult for the healthcare providers to address all needs (Wang et al., 2018). But to increase patient satisfaction and improve health outcomes, the complex care needs must be addressed without burdening the healthcare providers.

With significant technological advancements especially with the advent of Artificial Intelligence (AI) and Machine Learning (ML), there is an opportunity to efficiently bridge this societal gap while effectively addressing the multiple needs of these chronic patients. We see the application of AI-driven data science in areas like predictive analytics, personalized medicine, and natural language processing to improve diagnoses, treatments, and understanding of patient data (Anand, 2024).

This master thesis focuses on using data enabled design methods to personalize care pathways for chronic patients while also reducing the burden on care providers. We delve deep into using data science approaches to analyse large amounts of patient experience data to extract the underlying needs of patients which can then be used for personalizing care.

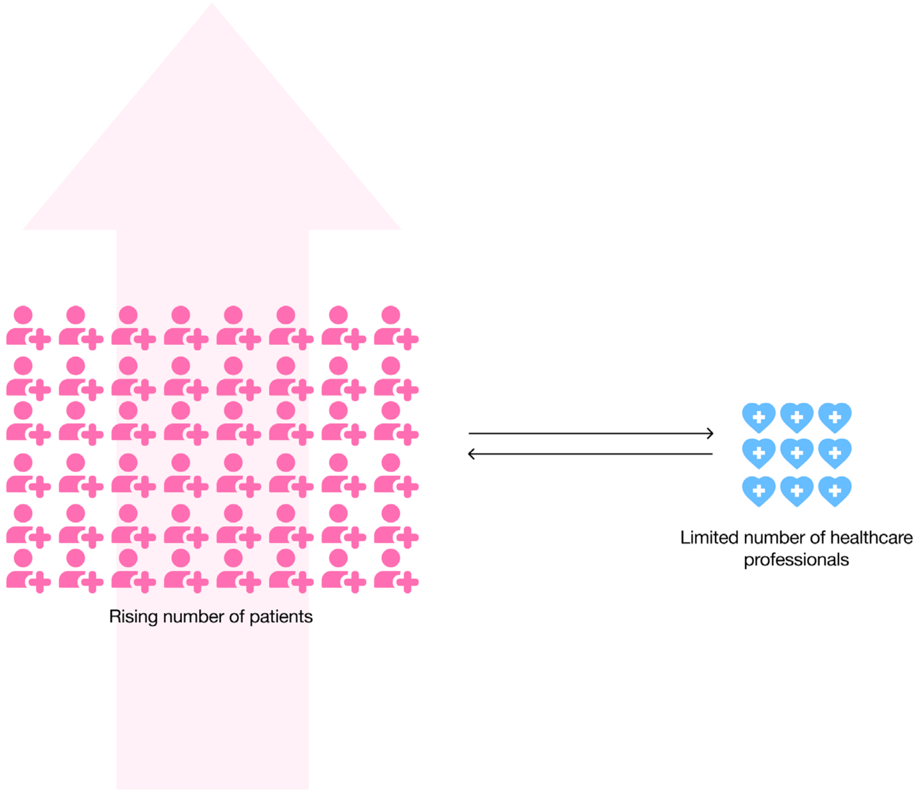

Fig. 1.1. Healthcare system under pressure

1.2 Central Research Question

To initiate the research conducted in this master’s graduation project, I began by identifying design opportunities based on outcomes from previous graduation projects (Kicken, 2023; Sun, 2023; Zhang, 2023; Quijada Fernández, 2024) and supporting literature. A recurring insight shared by clinicians was the need to help patients understand where they are in their care pathway.

A care pathway is an organizational tool aimed at improving patient care by promoting evidence, and guideline-based practices. It is a structured plan developed and used by multidisciplinary teams of doctors, specialists, nurse practitioners, and other healthcare professionals to guide various aspects of care, including assessment, diagnosis, investigation, and treatment (Kwan, 2007).

Care pathways offer several advantages when managing large patient populations: they standardize care, enhance communication among healthcare providers, reduce unnecessary procedures and medical errors, and maintain alignment with evidence-based practices (Kwan, 2007). However, despite their benefits, care pathways face challenges in their development, implementation, and evaluation. Moreover, creating these pathways is often resource-intensive, requiring significant time, effort, and dedicated personnel (Khalifa & Alswailem, 2015). These insights led to the central research question of this thesis.

> *“How can data-enabled design be used to create personalized care pathways that address the needs of chronic patients while reducing the burden on healthcare providers?”*

As I further carried out my research and design activities, the central research question was further broken down into a number of sub-questions that helped drive the necessary research and design activities. I will further introduce these research questions at the beginning of each chapter and describe the associated design activities along with outcomes and reflections from each chapter.

1.3 Stakeholders

This project is part of the Convergence Health & Technology flagship initiative, ‘Consultation Room of the Future’ at Erasmus MC, Rotterdam. It is a collaborative initiative that brings together researchers from Erasmus Medical Center, Erasmus University Rotterdam, and TU Delft focusing on tackling key issues affecting the accessibility and sustainability of healthcare, including workforce shortages, budget constraints, and evolving societal needs.

Fig. 1.2. Collaborating stakeholders

This project has a dual approach targeting mainly two user groups. The primary user group are healthcare providers, mainly clinicians of the three selected patient groups namely colorectal cancer, Pulmonary fibrosis and Sarcoidosis.The design outcome from this project will indirectly impact the associated patients making them the secondary, indirect stakeholders for this project.

1.4 Project Approach

The project approach is based on the enhanced data-enabled approach by Jung (2023). It consists of an iterative process where insights gained from data collection are continuously integrated into interpretation and design efforts.

The project also includes several data science based approaches to analysing large amounts of textual data making use of machine learning and AI based natural Language processing (NLP) techniques. In each chapter, I have further discussed the specific technique along with its outcomes in the form of quick experiments and reflection on the resulting outcomes.

In order to maintain reliability of the data being processed using machine learning (ML) or Large language models (LLMs), I have also incorporated specific verification steps from designers and clinical experts that are also discussed in the upcoming chapters. As a result, the research and design phases are closely connected, forming an ongoing feedback loop with primary stakeholders that guides the project’s development.

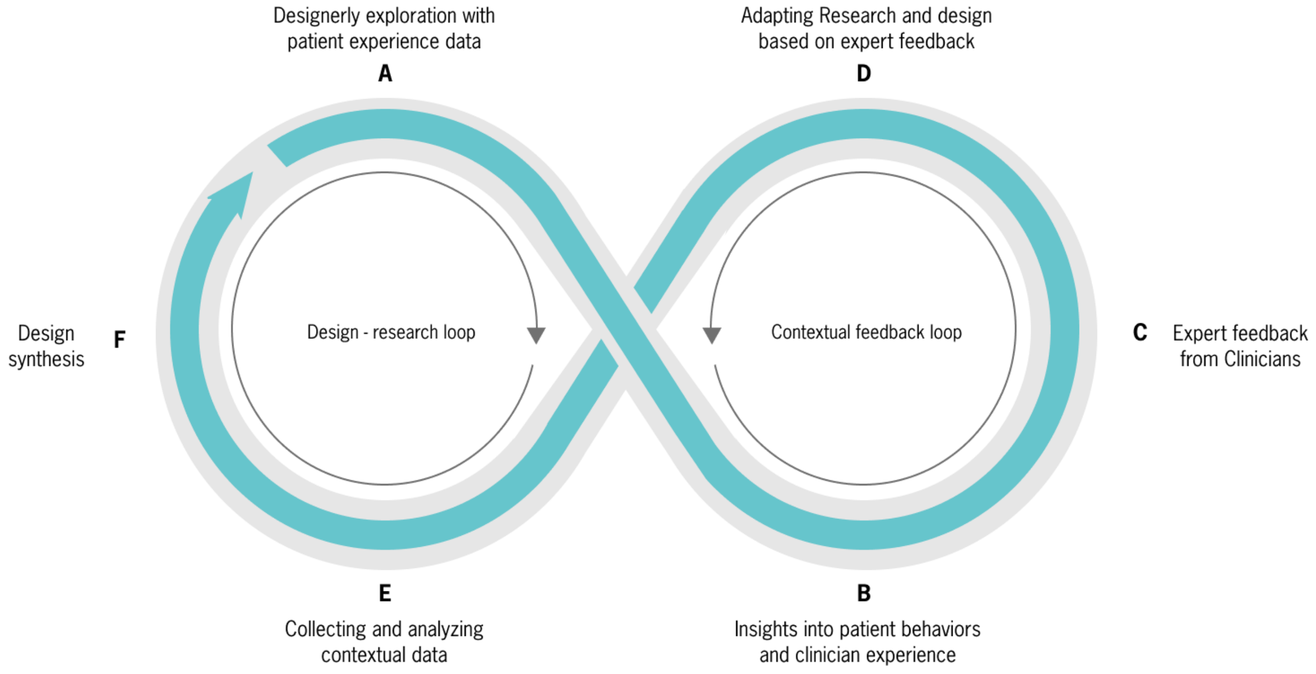

Fig. 1.3. Project approach visualization

1.5 Process Overview

This thesis followed a structured, yet deeply iterative research and design process aimed at translating large-scale, unstructured patient narratives into actionable insights for care personalization. The visual journey depicted in Fig 1.4. mirrors the layered methodology I developed, integrating data science, design thinking, and healthcare co-creation.

1. Reorganizing the Datasets

The journey began by using the structured patient datasets from previous topic modeling work (Quijada Fernández, 2024; Sun, 2023). This dataset consisted of disease-specific clusters, each tagged with human-interpreted topics and keywords. This set the foundation for NLP-driven interpretation using large language models.

2. Behavior Profiling with LLMs

I used GPT- 4o to extract representative patient behavior profiles. These profiles included goals, motivations, and challenges derived from keywords and cluster themes. I refined the LLM methodology through chain-of-verification (CoVe) prompting (Dhuliawala et al., 2024), ensuring contextual grounding and reducing hallucinations. Each behavioral insight was validated using the original patient blogs.

3. Manual Verification

Given the limitations of LLM verification, I introduced a manual verification phase, ensuring the personas reflected real patient experiences. Clinically significant discrepancies or generic patterns were flagged and corrected, resulting in 30 behaviorally rich personas across the three diseases.

4. Sorting Realistic vs Unrealistic Behaviors

With this rich persona library, I initiated expert co-creation sessions with healthcare professionals from Erasmus MC. Together, we sorted behaviors into realistic and unrealistic categories based on their clinical experiences, ensuring the behavioral models resonated with actual practice.

5. Mapping Behaviors on the Care Path

In a second co-creation activity, clinicians mapped each validated behavior onto different phases of the care pathway starting with prevention, diagnosis, treatment, and until follow-up and end-of-life care. This helped identify when specific behaviors tend to emerge, change, or overlap, and where gaps in support or misalignment occur.

6. Data Analysis and Thematic Synthesis

Behavioral data, clinician interviews, and sorting outputs were analysed to extract cross-disease themes, systemic challenges, and care delivery gaps. These insights revealed key areas of friction in care personalization, such as goal setting, patient literacy, and time constraints.

7. Identifying Opportunity Areas

The synthesis of co-creation sessions led to the articulation of shared unmet needs both from patients and clinicians. For example, patients needed tools that reduce emotional burden and improve health confidence, while clinicians needed behavioral insight dashboards that fit within their workflow.

8. Systemic Challenge Mapping

I constructed influence maps to visualize the broader clinical, organizational, and behavioral forces affecting care delivery. This helped surface structural problems, such as fragmented feedback loops and clinician dependence on intuition, which undermine personalization.

9. Strategic Vision for 2040

Based on these systemic insights, I defined a long-term strategic vision for care transformation. This roadmap, aligned with policy and technology forecasts (WHO, 2024; OECD, 2023), outlining three horizons - from piloting adaptive care tools in 2025 to embedding real-time behavioral AI into national systems by 2040.

10. Feedback Loop and MVP Design

Finally, I translated the strategy into a tangible service concept - a patient app and clinician dashboard. The app supports journaling, quality-of-life goal setting, community reflection, and PAM-based learning. The dashboard summarizes patient behavioral data, giving clinicians fast, behavioral insight before consultations. Further feedback from clinicians at Erasmus MC confirmed its potential for real-world impact.

This layered process turned scattered patient stories into a structured, scalable pipeline for care personalization. It bridges AI-driven insight with clinician intuition and culminates in a product that speaks to the emotional and systemic realities of modern care. In the upcoming chapter, we dive deep into the theoretical background to this study and discuss new approaches towards personalization of care.

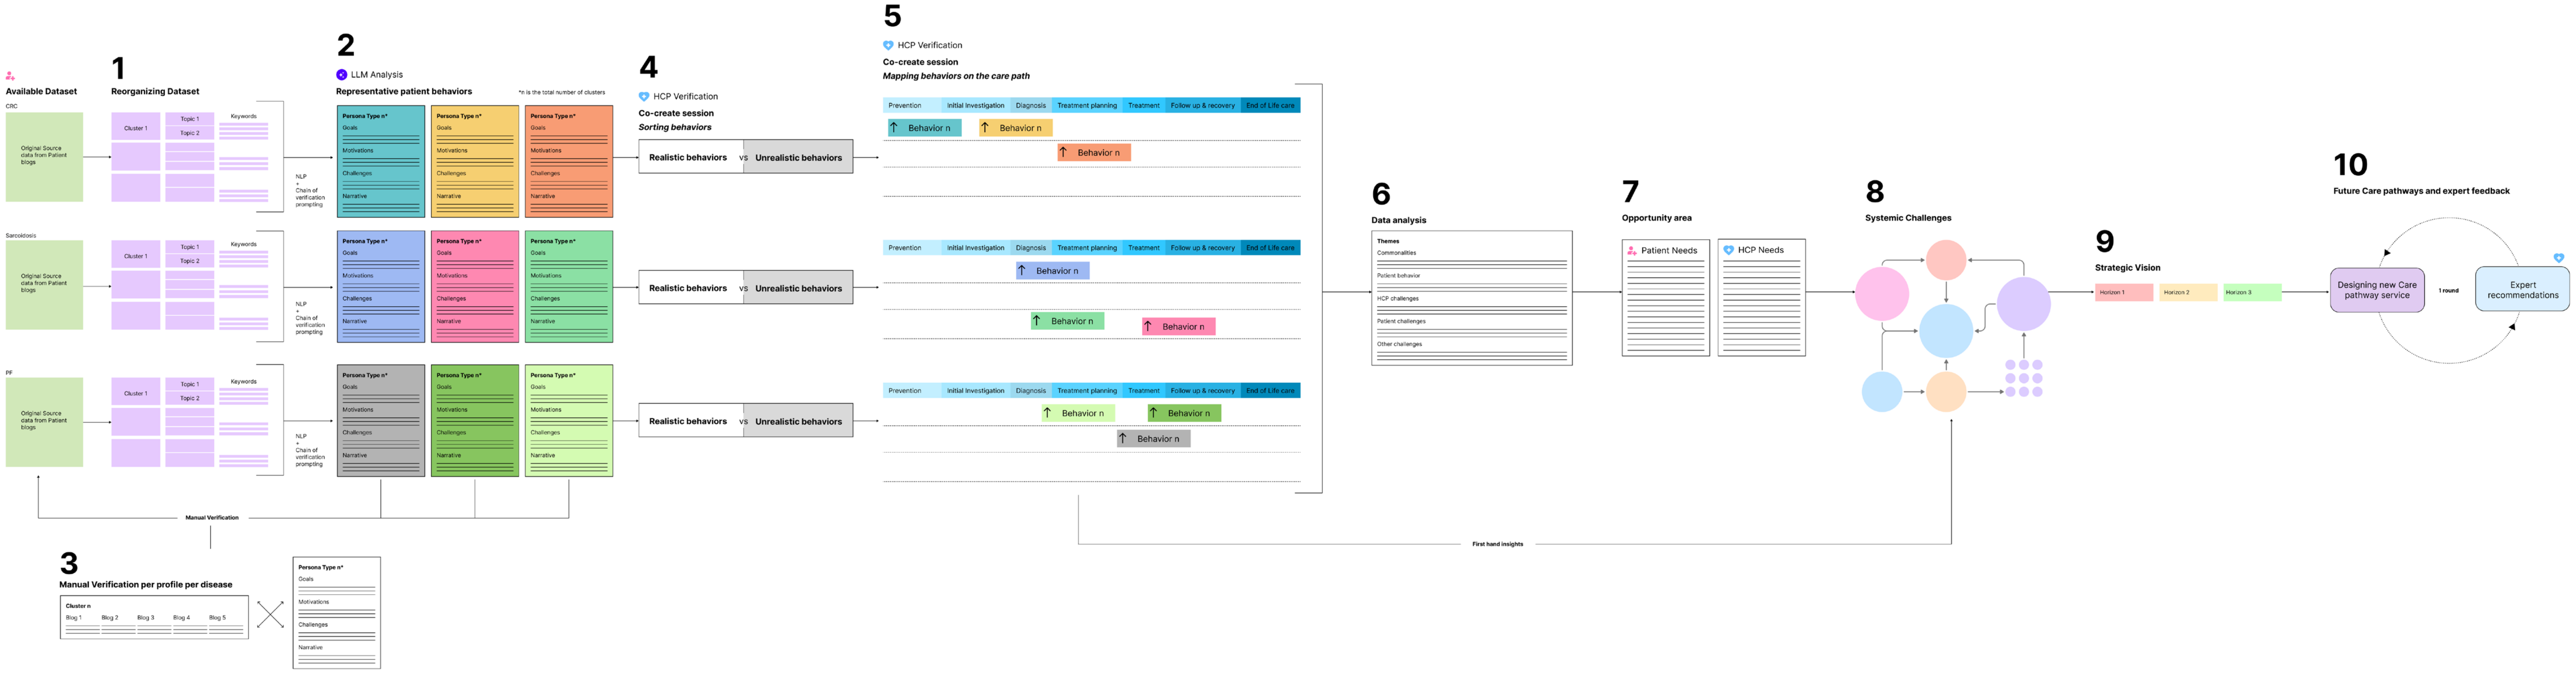

Fig. 1.4.Research and design activities of this project

Chapter 2
Working with data

This chapter explores the theoretical and methodological foundation for analyzing large-scale patient experience data. It addresses the central aim of leveraging these insights to design personalized care pathways. Since we are working with large amounts of structured data, we needed to tackle the following sub-research questions.

“How might we make use of previously collected and organized patient experience data?”

“How can clustered datasets be creatively utilized to create actionable deliverables that support personalization?”

“How do computational methods contribute to the accurate interpretation of patient needs?”

“What are the practical limitations in verifying automatically generated narratives from large language models?”

Patient experience data refer to written reflections shared by individuals in online patient communities. These narratives often appear on platforms that promote peer support. They form a rich and unstructured data source capturing the day-to-day realities, emotions, and uncertainties of people managing chronic or critical illnesses. This chapter outlines the processes used to organize, analyze, and repurpose this data, to inform the design of personalized care pathways.

2.1 Analysis of patient experience data

In earlier graduation projects (Kicken, 2023; Sun, 2023; Zhang, 2023), designers collaborated with data scientists to organize and analyze patient blogs from three patient communities including Sarcoidosis, pulmonary fibrosis and colorectal cancer.

The topics used in these previous projects were generated through a combination of computational methods and human interpretation. Using a topic modelling algorithm (Sushil et al., 2021), the data scientist grouped the patient posts into a predefined number of topics and identified key terms representing each of the topics.

After this step, the designer examined a sample of 50 posts per topic along with the extracted keywords to interpret and label the content accordingly (Jung et al., 2023). These labelled topics made it possible to compare recurring themes and lived experiences across different patient communities.

In another graduation project, designers further identified commonalities between the three different diseases by using computational analysis and verification with clinicians (Quijada Fernandez, 2024).

This resulted in a dataset that consisted of clusters of patient experiences, the associated topic under each cluster and related keywords per topic. This organized dataset served as a key source of information, driving several decision-making points throughout this project.

2.2 New strategy for personalization & first experiment

With the goal of personalizing care pathways for patients, it is imperative to properly understand the patient before offering support strategies. This would mean clearly understanding the patient experience and associated needs or concerns to create need based patient profiles. To identify the experience and needs from the clustered datasets, a six step plan (see fig 2.1) was proposed that made use of Natural Language processing through OpenAI LLM based text analysis.

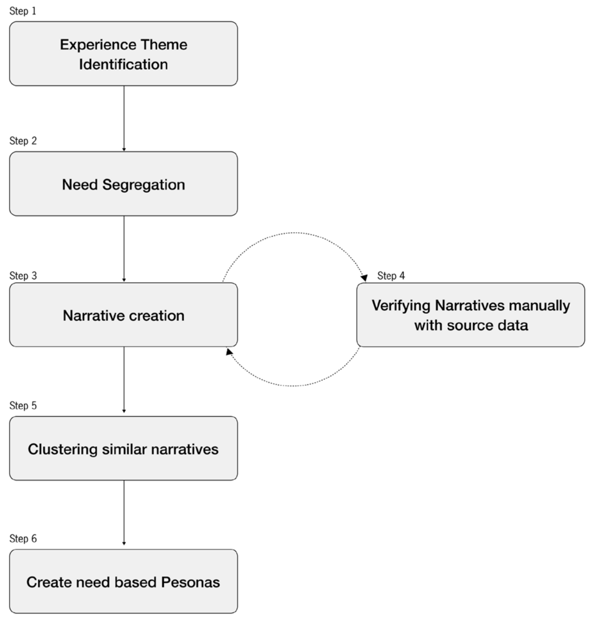

Fig. 2.1. Six step plan for creating patient profiles

Procedure

We decided to first identify experience themes from the patient data. The rationale behind this idea was to identify specific patient experiences and their associated behavioral patterns. To achieve this, we applied large language model prompting using the GPT-4o model to perform semantic clustering and thematic interpretation of textual data. We used clusters, topics, and keywords as input to generate experience themes that aligned with the original topics.

Next, we extracted patient needs from the experience themes and categorized them according to the five dimensions of health: physical, social, psychological, spiritual, and intellectual needs. After defining the needs and themes, we generated representative patient narratives using the cluster metadata and categorized needs.

Since the narratives were generated by a large language model, we implemented a verification step. This step involved tracing the generated narrative back to original patient posts within the same topic and validating its accuracy qualitatively.

We then applied unsupervised machine learning to cluster similar narratives across all disease categories. Finally, we used the verified themes, categorized needs, and clustered narratives to construct patient profiles that highlighted distinct goals, motivations, and challenges.

More about the data science method

The analysis in this study relies on transformer-based embeddings (text representations) developed using models like OpenAI's GPT-4o. These representations convert unstructured text into high-dimensional vectors that preserve semantic relationships between words, phrases, and entire passages. This method supports a range of advanced language understanding tasks by capturing both context and meaning within the data (Korade et al., 2024). The following points summarize the key advantages of using these representations in patient experience analysis.

1. Text Similarity: OpenAI embeddings can be used to measure the similarity between different pieces of text, which is essential for applications like document and textual clustering.

2. Sentiment Analysis: By capturing the nuances of language, embeddings improve the accuracy of sentiment analysis, helping businesses understand customer feedback and social media trends.

3. Named Entity Recognition (NER): OpenAI embeddings facilitate the identification of entities such as names, dates, and locations within text, which is essential for information extraction and knowledge graph construction.

2.3 First experiment results

To evaluate the proposed methodology, I used a small batch of clinician-verified data from the colorectal cancer community (Quijada Fernandez, 2024). This dataset served as the foundation for generating patient narratives, extracting needs, and testing the overall workflow in a secure and controlled environment.

There were a total of 10 clusters in the CRC dataset with a total of 38 topics with top 20 keywords for each topic (see fig 2.2 for example of dataset). This was used as the source dataset for extracting the experience themes, needs and narratives.

A simple prompt was used by providing the role of an expert healthcare provider to the LLM along with the input dataset (see fig 2.3). To avoid hallucination and unnecessary generative content, the goals for extracting the associated need was clear and specific. Each need was individually defined and also relied on the role of the LLM as an expert to be able to identify these needs.

Cluster	Topic	Keyword
Cluster 1: Experience around medical professionals's opinions	Doubts about treatment opinions from medical professionals	['doctor', 'oncologist', 'surgeon', 'told', 'patient', 'system', 'cell', 'wrong', 'reason', 'medical', 'asked', 'office', 'immune', 'trust', 'kill', 'testing', 'medicine', 'appointment', 'clinic', 'recommend']
	Suggestion to look for a second medical opinion	['care', 'second', 'opinion', 'ask', 'port', 'onc', 'put', 'taken', 'nurse', 'center', 'forget', 'top', 'first', 'team', 'comfortable', 'third', 'schedule', 'getting', 'question', 'hospital']

Fig 2.2. Example of one row of input data

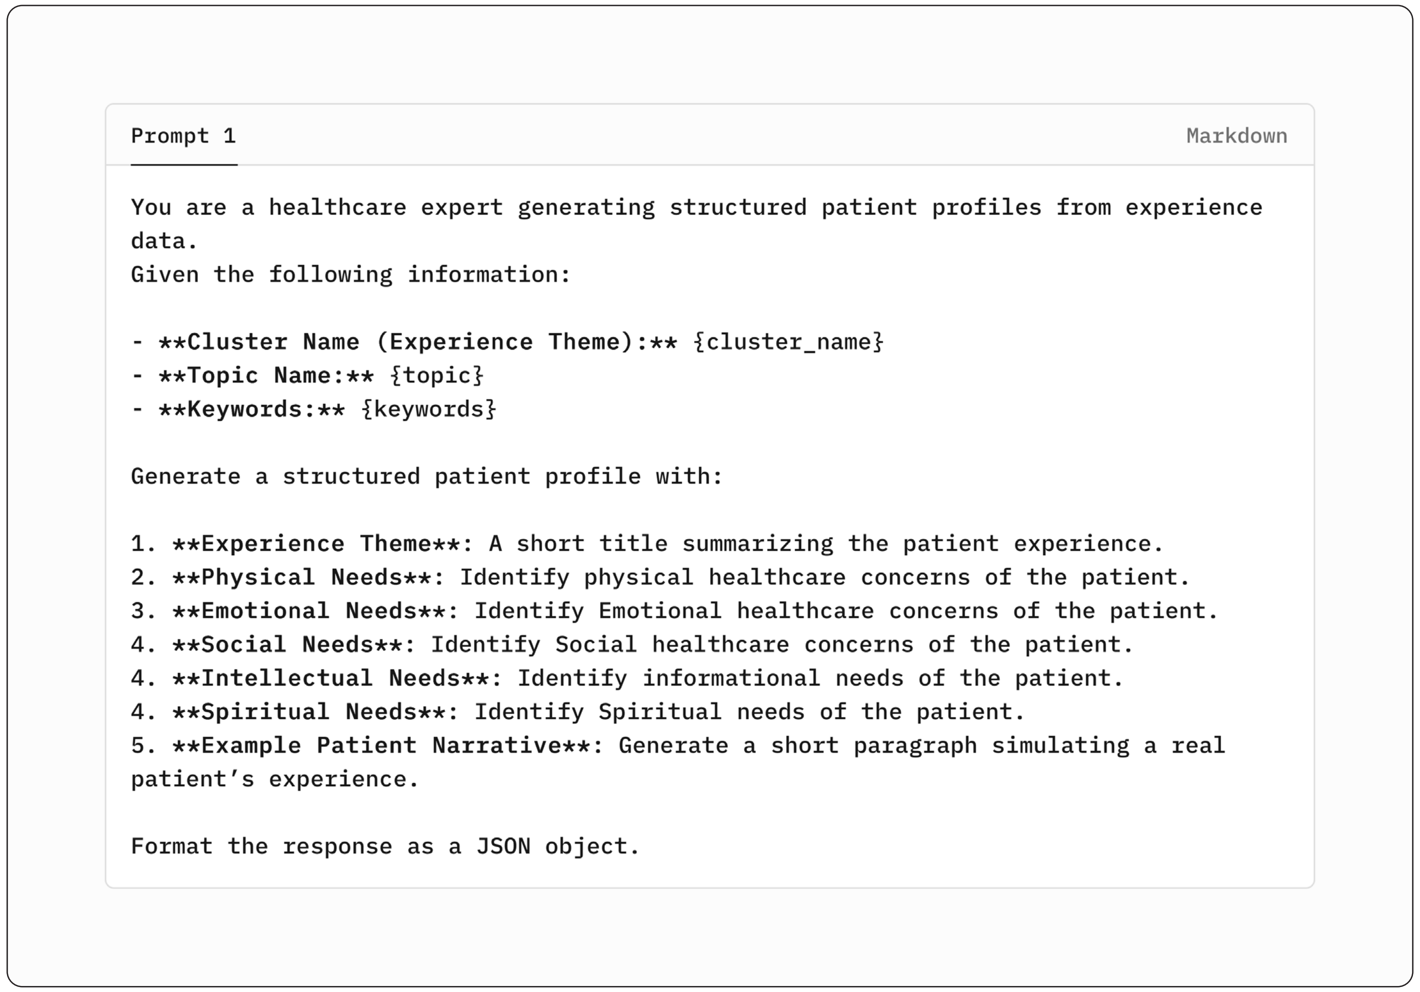

```
You are a healthcare expert generating structured patient profiles from experience
data.
Given the following information:

- **Cluster Name (Experience Theme):** {cluster_name}
- **Topic Name:** {topic}
- **Keywords:** {keywords}

Generate a structured patient profile with:

1. **Experience Theme**: A short title summarizing the patient experience.
2. **Physical Needs**: Identify physical healthcare concerns of the patient.
3. **Emotional Needs**: Identify Emotional healthcare concerns of the patient.
4. **Social Needs**: Identify Social healthcare concerns of the patient.
4. **Intellectual Needs**: Identify informational needs of the patient.
4. **Spiritual Needs**: Identify Spiritual needs of the patient.
5. **Example Patient Narrative**: Generate a short paragraph simulating a real
patient's experience.

Format the response as a JSON object.
```

Fig 2.3. LLM prompt used for extracting needs

Results

We were able to extract the associated needs and generate the representative narratives topic wise. 38 different narratives were created with each topic having its own set of physical, emotional, social, intellectual and spiritual needs.

Cluster Name	Experience Theme	Physical needs	Emotional needs	Social needs	Intellectual needs	Spiritual needs	Sample Patient Narrative
Cluster 1: Experience around medical professionals opinion	The patient's experience revolves around seeking multiple medical opinions, with a focus on oncology. They are navigating the healthcare system, scheduling appointments, and interacting with various medical professionals.	Primary Concern: The patient has a serious health condition, potentially cancer, requiring the involvement of an oncology team. Secondary Concerns: The need for a port to be put in for treatment, The physical strain of attending multiple appointments	Primary Emotion: Uncertainty Secondary Emotions: Anxiety about the diagnosis and treatment, Discomfort with the first medical opinion, Stress from managing appointments and seeking further opinions	Primary Need: Support from family and friends during the process of seeking multiple opinions and managing healthcare decisions. Secondary Need: Understanding and empathy from medical professionals.	Primary Need: Clear, comprehensive information about their diagnosis and treatment options. Secondary Needs: Guidance on how to seek a second or third opinion, Understanding the process and implications of getting a port put in	Primary Need: Reassurance and hope during a challenging time, which may come from personal beliefs, faith, or a sense of purpose.	The patient, facing a potential cancer diagnosis, is navigating a complex healthcare journey. They have been advised to have a port put in for treatment, but are uncomfortable with the first opinion received. They are now seeking second and even third opinions, scheduling appointments with different oncology teams. The process is physically and emotionally taxing, causing anxiety and uncertainty. They rely on their close social network for support, and seek clear, comprehensive information to make informed decisions. Their spiritual beliefs provide them a source of hope and resilience during this challenging time.

Fig 2.4. Example of one extracted experience theme with categorized needs and narratives.

Verification with original blog post

Each need and the generated patient narrative was verified with the original patient blog from the related topic. For example, the data from fig 2.4, was verified by reading and cross checking 3-5 original patient blogs picked randomly from the top 50 posts of the topic “Doubts about treatment opinions from medical professionals” which was under the cluster of “Experience around medical professional’s opinion” (see example of original patient blog on the same topic on the right). This method helped us check for hallucinations in the generated content.

Correlating the original blog post qualitatively with the data from fig 2.4, the generated data was accurate in describing the needs and the narratives of the patient. This is mainly because the experience themes, needs and narratives were extracted from a well defined topic cluster with appropriate keywords.

The verification procedure consisted of individually verifying the generated needs of the patient with at least 3 individual blog posts. However, this verification process in itself was labor intensive and took significant time to cross verify each column of the generated dataset with the original patient blogs.

Original blog post

> *Hi, I am kind of in the middle of this treatment. I still have my temporary ileostomy.. I had most of my rectum removed and several inches of my descending colon. Then they pulled down the end of the colon and made a J Pouch to serve as my new rectum. They also saved my sphincter of course. I have around 2/3 of my colon remaining. I am under the impression that the more colon that remains intact, the better chance of returning to normal bowel movements. I will have my ileostomy reversed at the beginning of the year. My surgeons expect me to return to normal. We will see though! I am also going to be having surgery to fix some adhesions in my small intestines which is scar tissue that can kink and cause blockages (not fun). So after the first of the year I will be in for a ride to see how well my guts are behaving! Cherry on... my gall bladder will be removed (they are also removing 1/3 of my liver with a met in it) 3-for-one surgery! That should be interesting living without a gall bladder... I should post about what that is like... although I have been told it’s not really noticeable. Anyhoo... I am ramblin’... upshot is my 2 cents is yet to come on your question... I am hoping that I can believe my surgeons when they speak of my returning to normalcy.*

2.4 Reflections and refinement

Upon reflecting on this methodology, extracting needs and creating accurate patient narratives were easy, but there were three significant challenges that had to be addressed.

Difficulty in defining the needs
The needs defined in the LLM prompt were not based on specific theories. To accurately define the needs, I explored different theory-based models starting with the 5 dimensions of health (J. Hjelm, 2010), adapted version of Maslow's hierarchy of needs (Carroll et al., 2025), Chronic care model (Bodenheimer et al., 2002), Quality of Life (EURO QoL+CRC) questionnaires, systematic reviews (Wang et al., 2018) and other qualitative studies. The definition of 'needs' varied significantly from one model to another. Either the needs described were too extensive or too concise (López-Salas et al., 2024) and always overlapped with each other (with one need influencing another) making it significantly difficult to extract in isolation. Moreover the extracted needs were generic for the experience theme making it difficult to cluster and create personas.

Flaws in experience theme
The experience theme almost always ended up being a paraphrase of the cluster or associated topic making the generated data redundant. It did not add value to the generated dataset.

Feasibility issues
Since the needs were extracted and narratives were generated on a per-topic basis, the extracted needs or patient concerns (physical, emotional, social, intellectual, and spiritual) had to be individually verified against the original blog posts, making the process labor-intensive and time-consuming. This involved verifying five needs per topic for approximately 150 topics across the three selected diseases. Although the generated narratives were accurate representations of the patients, they shared the same issue of being labor-intensive to verify individually.

2.5 Conclusions

The experiment demonstrated that language models can effectively transform unstructured patient narratives into structured profiles when given well-defined prompts and high-quality input. The approach holds promise for supporting personalized health interventions.

To improve the methodology, future work should focus on refining need taxonomies, automating partial verification using similarity metrics, and enhancing the LLM prompt structure to reduce redundancy. Addressing these areas can make the process more scalable and clinically actionable.

Chapter 3 Creating patient profiles

Since defining broader health needs for accurate extraction of patient needs was a challenge, this chapter builds on the previous experiment by refining the approach to generating patient profiles. Instead of extracting broad health-related needs, I focused on identifying goals, motivations, and challenges as key behavioral indicators. This new approach introduced technical challenges, such as hallucinations and data accuracy issues, that needed to be addressed leading to the following sub-research questions.

“How might we reduce hallucinations and improve the accuracy of the generated content?”

“How might we ensure the generated personas are contextually accurate?”

“How might we optimize the content verification procedure for feasibility and scalability?”

This chapter outlines the revised methodology, presents generated personas, and reflects on key improvements and remaining limitations.

3.1 Refinements in methodology

Based on the learnings from the previous experiment, the methodology was modified to avoid extracting needs of the patient that were difficult to define. Instead, text-based analysis was used to directly generate the patient profiles.

Goals	An aim or the desired result of the patient
Motivations	Reason(s) for the patient acting or behaving in a particular way
Challenges	Obstacle(s) preventing the patient from achieving their goals

Table 3.1. Defining the key terms used in LLM prompt

Instead of trying to define the needs of the patient, we decided to identify the goals, motivations and challenges of the patient. Goals, motivations and challenges can be defined objectively and provide a clear picture of the associated patient behavior.

3.2 Optimizing the verification process

When using large language models to generate patient data from topic clusters, we observed several instances where the output included irrelevant, generic, or factually inaccurate content. These errors, commonly referred to as hallucinations, occurred when the model produced statements that did not align with the original dataset or failed to reflect the intended patient behavior.

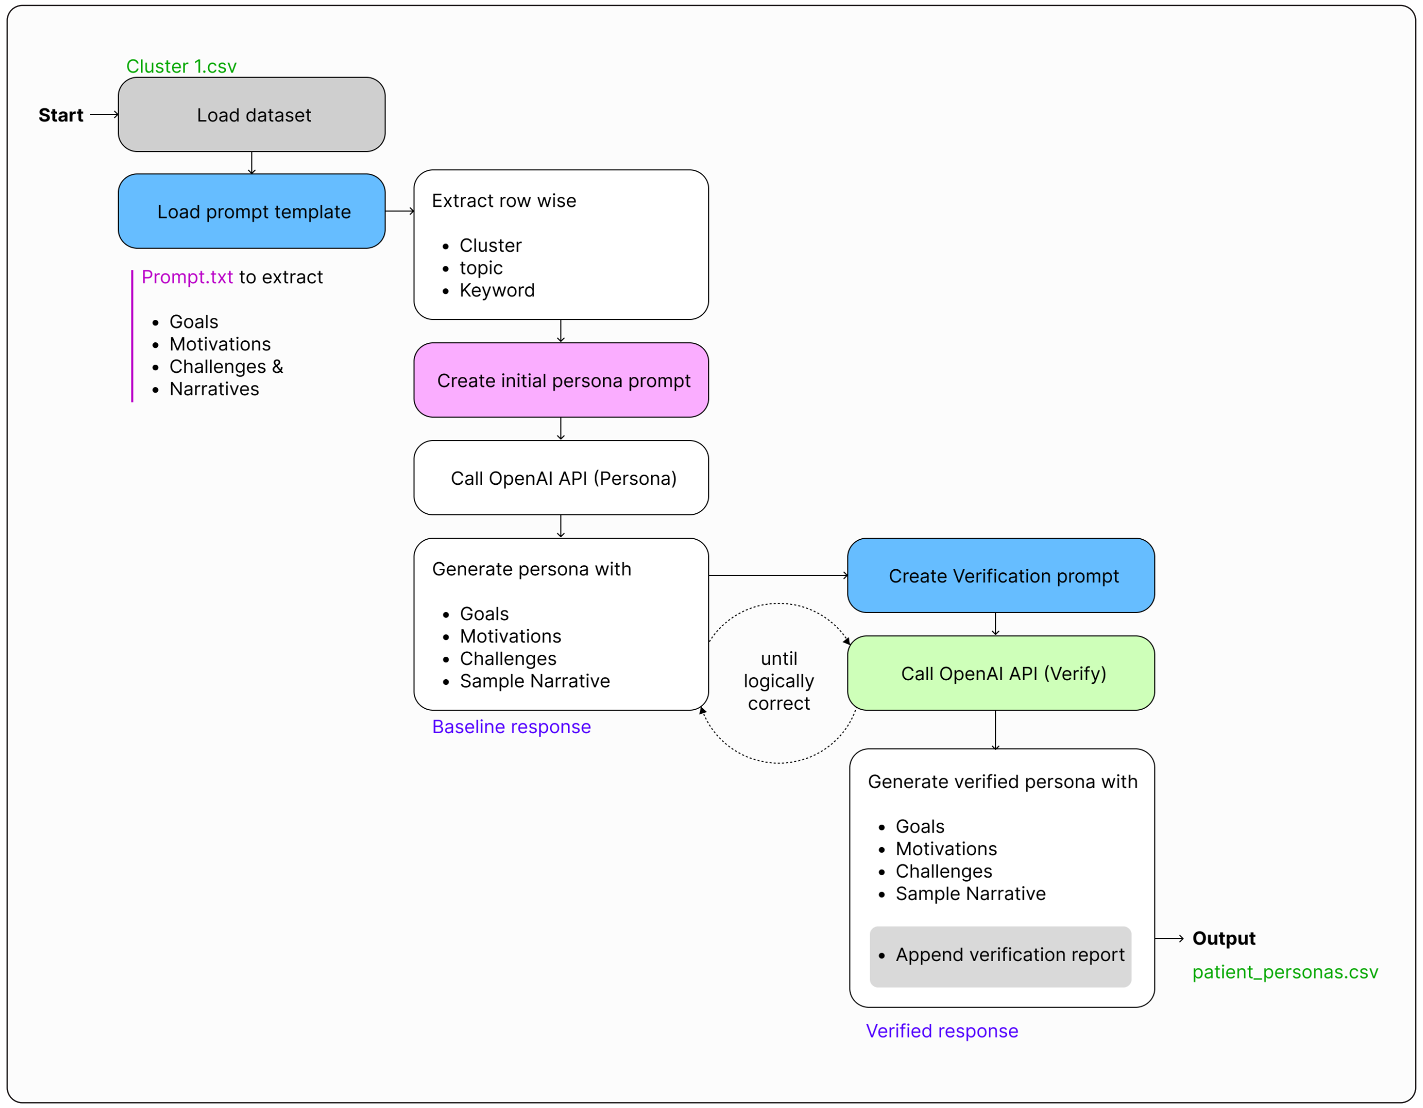

Fig 3.2 Processing logic flowchart

To reduce these hallucinations, we implemented a verification strategy that allowed the model to evaluate and refine its own output. This method, known as Chain of Verification (CoVe) prompting, involved generating an initial response, creating targeted verification questions, answering those questions independently, and revising the original content based on the findings. This multi-step approach helped us make the model reason more logically and improved the overall accuracy of the generated personas (Dhuliawala et al., 2024).

The processing logic was modified to incorporate this methodological step (see fig 3.2. flowchart below). Once the datasets are loaded, the clusters, topics and keywords are extracted row wise. This information is then used to create a base prompt for creating the persona. Once the OpenAI API is called, a baseline response is generated which contains the first versions of the generated persona that includes goals, motivations, challenges & narratives.

After this step, a verification prompt is generated with the objective of the LLM to verify its own response. The Open AI API is called for the verification prompt, and the baseline response is evaluated until it makes logical sense. The final verified response consists of the original persona with the appended verification report. The final persona document is exported as a new CSV file.

3.3 Generating patient profiles

Using this new methodology, patient personas were generated for a batch of CRC datasets that consisted of 10 clusters, 38 topics and 20 keywords per topic. Personas were generated cluster-wise for feasibility reasons. An example persona that was generated for the cluster "Experience around medical professional's opinions" is shown on the right (see fig 3.3) along with the original verification report (see fig 3.4).

1. Type of Persona: The Doubtful Patient

2. Goals of Persona:
Seeking clarity about their treatment: This persona needs to understand what the doctor, oncologist, or surgeon has told them about their medical condition.
Getting a second opinion: They believe in the importance of having a second or even a third opinion to ensure the accuracy of the diagnosis and treatment plan.
Ensuring the best care possible: They want to be comfortable with their medical care team and the recommendations made for their treatment.

3. Motivations:
Trust in the medical system: They have a strong desire to trust the medical professionals (doctor, oncologist, surgeon, nurse) and the healthcare system.
Ensuring a thorough understanding: They are motivated to ask questions and get detailed information about their condition and the suggested treatment.
Desire for the best possible outcome: The persona is motivated to get the best medical care and treatment to improve their health condition.

4. Challenges:
Lack of trust: They sometimes have a feeling that something is wrong with the opinion given by the medical professional.
Communication barriers: They may have difficulty understanding complex medical terms and procedures that the doctor or oncologist explains.
Fear and uncertainty: They may feel overwhelmed and anxious about their health condition, the treatment options, and the potential outcomes.

5. Narrative:
This persona represents a patient who has been diagnosed with a serious health condition. They have been given a treatment plan by their doctor or oncologist, but they are struggling to fully trust this opinion. They find the medical jargon confusing and fear that they might not be getting the best care possible. This leads them to seek a second or even a third opinion. They are motivated by their desire for clarity and the best possible outcome, but they are challenged by their doubts and fears. They are willing to ask questions and seek further information to become comfortable with their care team and the suggested treatment plan.

Fig 3.3. Generated and verified persona

Verification Report

The patient persona, "The Doubtful Patient", is generally accurate and complete. It provides detailed information on the persona's goals, motivations, challenges, and narrative. The persona is well-defined, relatable, *and realistic, representing a significant portion of patients who are dealing with serious health conditions and are unsure about their treatment.*

However, it might be beneficial to add more specific details to make the persona more realistic and relatable. For example, adding demographic information such as age, gender, or occupation could provide a more concrete picture of the persona. Also, the persona could be made more relatable by including personal details such as hobbies or family life.

The narrative section could also include how this persona typically interacts with healthcare providers and how they gather additional health information (e.g., online research, medical literature, discussion with family or friends).

Overall, the persona is well-structured and provides a good insight into the mindset of a doubtful patient. With some minor improvements, it could serve as a valuable tool for healthcare providers to understand and address the concerns of such patients.

Fig 3.4. Verification report for associated persona

3.4 Issues with generated patient profiles

Although the example of the generated persona is complete, some of the other generated personas for other clusters consisted of incomplete narratives (see fig 3.5 below). Some personas were generated without referencing the keywords that were used. Inconsistencies were also observed in the formatting of the generated data. Another observation from studying the personas and the verification report, is that the verification report is generic and suggestive.

Although the LLM has analyzed its own response, there is no real verification on the content that is being generated. The verification process did not correct the actual response, instead, it provided a verification report. A manual verification step was still necessary in order to verify the accuracy of the content generated.

To do this, I revisited relevant topics from the original patient blogs and randomly selected 3 blogs and cross checked it with the personas generated. The generated personas were closely representative of the original patient blogs but there were still some discrepancies and generic data that lacked specificity.

1. Type of Persona: Proactive Health Researcher

2. Goals of Persona:
Conduct comprehensive research on alternative therapies (Keywords: ‘research’, ‘alternative’, ‘therapy’, ‘study’, ‘internet’, ‘article’)
Make informed decisions about future treatments (Keywords: ‘decision’, ‘future’, ‘choice’, ‘outcome’, ‘decide’, ‘path’)
Understand and manage drug side effects and regimen (Keywords: ‘side’, ‘effect’, ‘drug’, ‘regimen’, ‘dose’, ‘reaction’)
Explore the use of Traditional Chinese Medicine and supplements (Keywords: ‘chinese’, ‘managing’, ‘supplement’, ‘medication’)
Stay informed about clinical trials (Keywords: ‘clinical’, ‘trial’, ‘finding’)

3. Motivations:
Desire to achieve the best possible health outcome (Keywords: ‘best’, ‘health’, ‘outcome’)
Interest in alternative and adjunctive therapies (Keywords: ‘alternative’, ‘therapy’, ‘consider’, ‘option’, ‘approach’)
Willingness to share and receive information with/from others (Keywords: ‘information’, ‘share’, ‘experience’, ‘suggestion’)
Belief in making informed decisions (Keywords: ‘decision’, ‘information’, ‘choice’, ‘decide’)

4. Challenges:
Difficulty in finding reliable sources of information (Keywords: ‘information’, ‘research’, ‘website’, ‘article’, ‘internet’)
Managing the side effects of drugs (Keywords: ‘side’, ‘effect’, ‘drug’, ‘reaction’, ‘manage’)
Understanding the effectiveness of different treatments (Keywords: ‘effective’, ‘treatment’, ‘responded’)
Navigating the complexities of clinical trials (Keywords: ‘clinical’, ‘trial’, ‘complexity’, ‘navigate’)

5. Short and Concise Narrative:
This persona represents a patient who is highly proactive in managing their health. They are diligent in researching alternative therapy options and are keen on making informed decisions for future treatments based on the best possible outcomes. They are interested in understanding the side effects and regimen of drugs, and are open to exploring Traditional Chinese Medicine and supplements. They appreciate the value of sharing experiences and recommendations, and are always on the lookout for new clinical trials. However, they face challenges in finding reliable information, managing drug side effects, understanding treatment effectiveness, and navigating the complexities of clinical trials. Despite these challenges, their motivation to

Fig 3.5. Narratives are truncated and in inconsistent format

3.5 Reflections and refinement

The shift from need-based profiling to behavior-oriented persona generation significantly improved the structure and interpretability of the output. By focusing on goals, motivations, and challenges, I was able to extract behavioral traits that felt more clinically usable and design relevant. These elements were easier to define objectively and allowed for clearer differentiation between persona types.

The use of Chain of Verification (CoVe) prompting played a key role in reducing hallucinations. The model became better at maintaining internal consistency, and the added verification layer helped identify logical gaps in its own responses. However, the CoVe process mostly served as a diagnostic tool rather than a corrective one. The verification reports often suggested improvements but did not revise the generated personas directly. As a result, manual review remained essential to ensure alignment with patient blogs and to correct vague or repetitive phrasing.

Despite these improvements in methodology, the overall quality of the personas varied across clusters. Some profiles were realistic and detailed, while others lacked specificity or failed to integrate key keywords from the topic data. Incomplete narratives and inconsistent formatting were common when the number of tokens approached model limits. These issues reduced the emotional fidelity of some personas and highlighted the need for further optimization in prompt structure and data sampling.

The process became more efficient and less prone to hallucinations, but this came with a tradeoff in creative depth. Many outputs remained safe, generic, or overly cautious. This revealed a limitation in the model’s ability to generate behavior-rich content without drifting into abstraction. It also underscored the importance of grounding personas in real patient language to preserve authenticity.

Moving forward, future iterations must focus on improving narrative richness without sacrificing reliability. This could involve integrating few-shot examples from original patient blogs or refining prompts to encourage more expressive storytelling. Verification strategies should evolve to assess content and revise it, ideally with human input in the loop.

Overall, this phase of the project succeeded in stabilizing the generation process but showed that behavioral nuance and contextual realism still require further development. The results provided a solid foundation, but they must be enriched further before these personas can fully support clinical or design applications.

3.6 Conclusions

This chapter presented a refined methodology for generating patient profiles using natural language processing and large language models. By focusing on behavior-based attributes specifically goals, motivations, and challenges, I moved away from abstract need classifications and towards a more structured and actionable persona format.

The introduction of Chain of Verification (CoVe) prompting helped reduce hallucinations and improved internal logic within the model’s outputs. However, while the process increased consistency, it did not fully eliminate the need for manual validation. Verification reports identified issues but did not correct them, leaving gaps in completeness and contextual grounding.

The results showed that while the generated personas were more structured and interpretable, they sometimes lacked specificity, emotional detail, or strong alignment with the original patient narratives. These limitations became most visible in clusters where blog data was sparse or where the prompts failed to elicit expressive responses from the model.

In conclusion, the revised approach improved reliability and reduced error, but it did not yet reach the level of contextual richness required for real-world clinical integration. Manual oversight remained necessary to ensure quality, and future iterations must aim to balance accuracy with emotional and behavioral depth. These lessons informed the next phase of the project, where the personas undergo validation and contextual enrichment.

Chapter 4
Creating Contextual personas

From the results of the previous experiments, the issues of hallucination had been addressed, but an optimization of the methodology was required to make the personas contextually accurate and emotionally rich. This resulted in the following research questions

“How might we ensure the generated personas are contextually accurate?”

“How might we optimize the content verification procedure for feasibility and scalability?”

This section outlines the computational pipeline designed to generate validated patient personas from a large dataset of patient blogs. The pipeline leverages structured prompting, and an iterative LLM-based verification loop to ensure the generation of high-quality and contextually grounded behavioral characteristics of patients. The process is implemented using OpenAI’s GPT-4o model and adheres to principles of prompt engineering and human-centered validation.

4.1 New Computational pipeline

To ensure the extraction and generation of contextually grounded personas, significant contextual data was required to create a knowledge base for the LLM. This knowledge base would act as a starting point for the generation of specific characteristics including patient goals, challenges and motivations painting a rich picture of the associated patient behaviors.

The proposed computational model follows a structured, multi-step process designed to ensure the accurate generation and verification of thematic patient personas. Refer to the workflow diagram in fig 4.1.

4.2 Computational Workflow

The computational workflow proceeds in the following steps:

1. **Initialization and loading datasets**
   The process begins with the initialization of the computational environment and the loading of two primary datasets for analysis. The first dataset is unstructured and comprises the top 50 original patient blog posts per topic, identified using predefined topic IDs. The second dataset is structured and includes a hierarchical organization of data, consisting of cluster names, associated topic names with their corresponding topic IDs, and keywords relevant to each topic. In this hierarchical framework, each cluster contains multiple topics, which are identified by their unique topic IDs, and each topic is characterized by a set of associated keywords. These clusters and topics were derived from prior research that applied topic modeling techniques to convert unstructured patient narratives into a structured format of clusters, topics, and keywords.

2. **Iterative cluster wise processing**
   Each thematic cluster is processed individually in a sequential manner to preserve contextual integrity and ensure focused analysis. In each iteration, only one cluster is considered, along with its related data such as topic IDs, topic names, and associated keywords. This approach helps maintain thematic consistency and allows the following steps to operate within a clearly defined context

3. **Knowledge base construction**
   Each cluster is processed by explicitly matching the topic IDs between the unstructured and structured datasets. The matched topic IDs, along with their associated topic names, keywords, and patient blogs, are temporarily stored. Keywords from the structured dataset are then used to further filter the original unstructured dataset. From this filtered set, five posts, representing approximately ten percent of the original blogs, are sampled and stored to form a knowledge base that provides contextual information for the language model. This step is critical to ensure that the model generates contextually accurate outputs.

4. **Verification Context creation**
   A verification context is created for each cluster individually. This is achieved by matching the topic IDs between the structured and unstructured datasets. Once the topic IDs are aligned, the language model randomly samples five original patient posts within each topic. The sampling logic is designed to exclude any posts that were previously selected for the knowledge base. This approach helps prevent complete overlap between the knowledge base and the verification set, thereby reducing the risk of bias and ensuring a more robust and reliable verification process.

5. **Prompt Construction**
   Prompts are constructed using a hybrid strategy that integrates few-shot prompting, instructional prompting, and chain of verification prompting. This approach combines curated examples and reference datasets to guide the generation of patient personas. Posts extracted for the knowledge base serve as few-shot examples, offering contextually relevant information to the language model.

Instructional prompting is employed to explicitly generate key persona elements, including patient goals, motivations, challenges, and a representative narrative. These components are not predefined within the prompt; rather, they are framed as a task requiring the model to infer and list characteristics directly related to the given cluster, topics, topic IDs, and associated keywords.

This strategy is essential to reduce hallucinations, ensure that the generated content remains grounded in the original data, and promote traceability and transparency in the persona development process. The entire prompt used is shown in fig 4.2 in the following page.

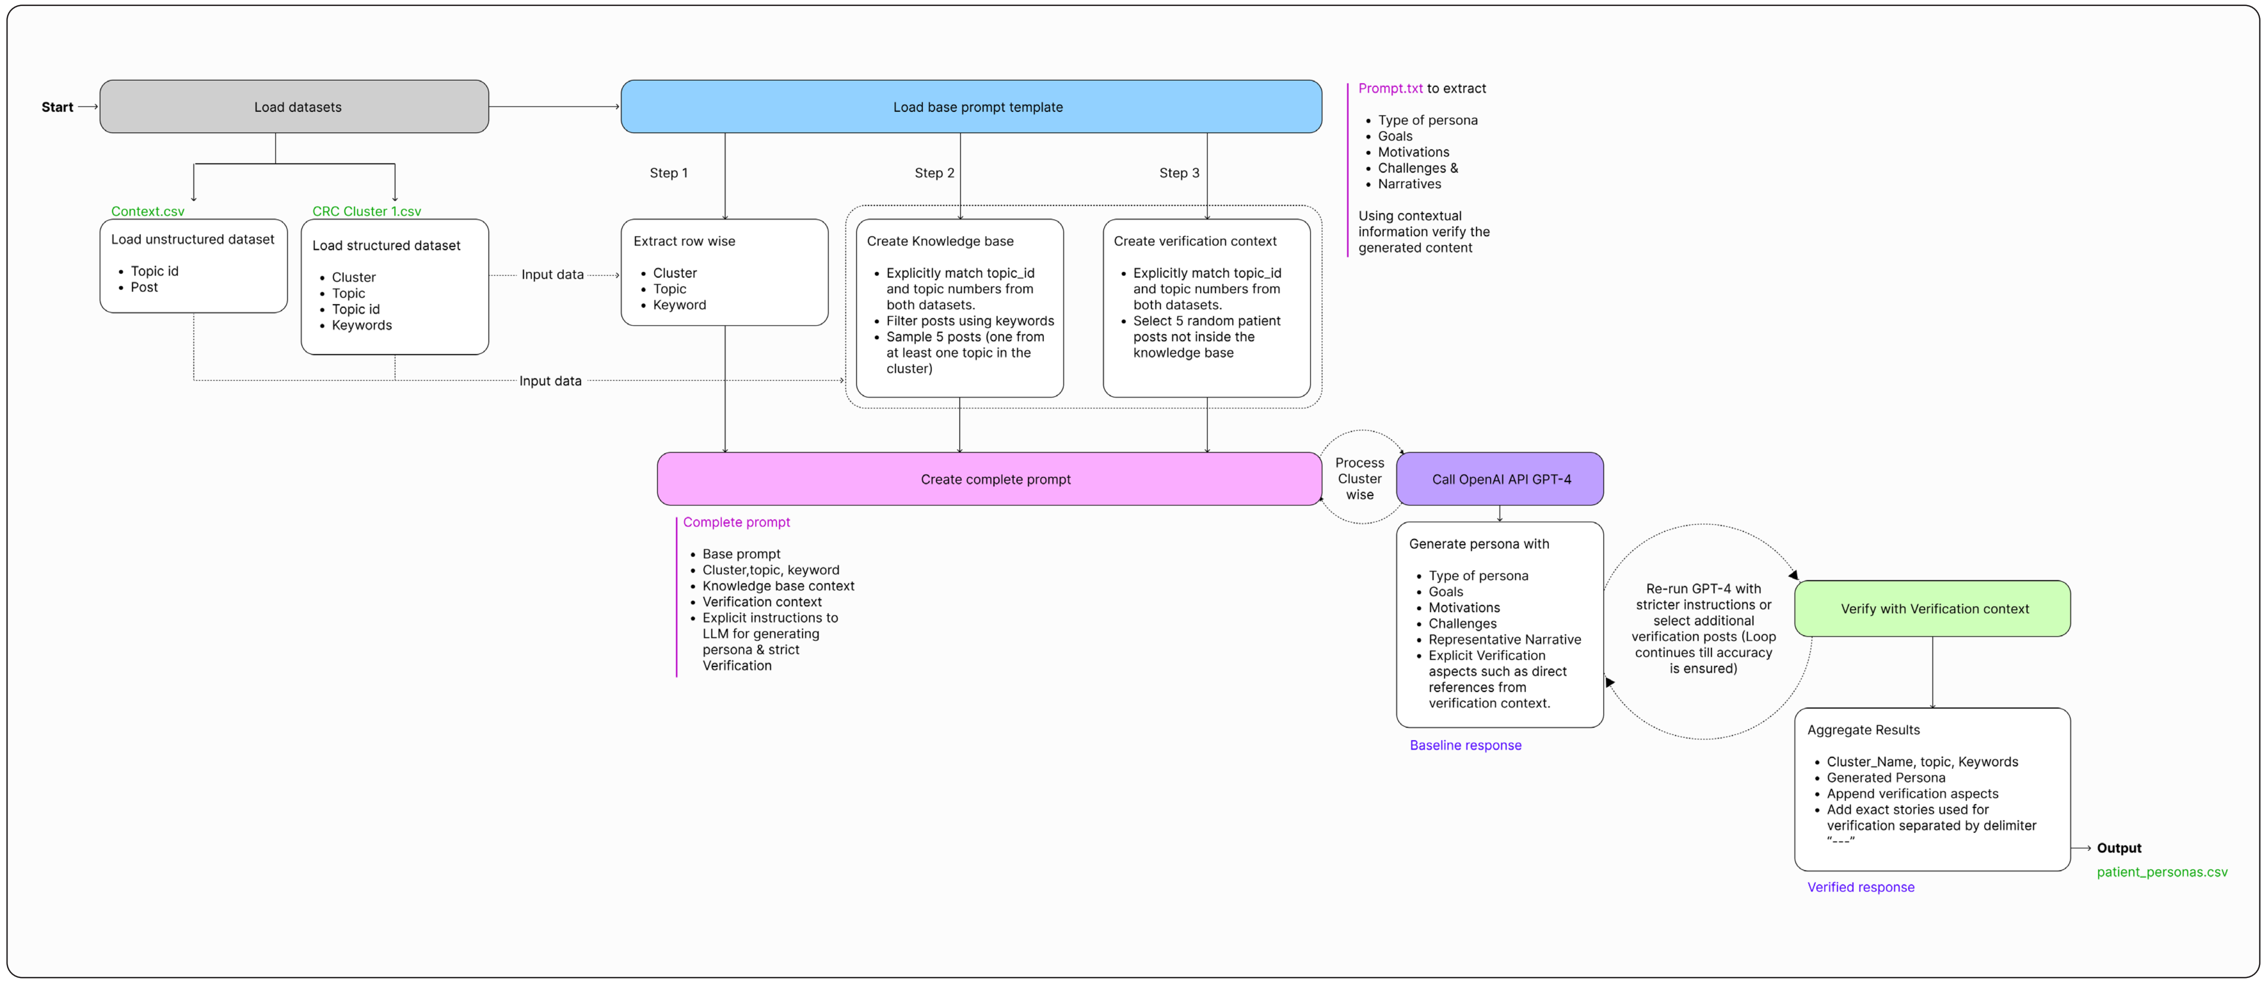

Fig 4.1. Workflow diagram of proposed computational model

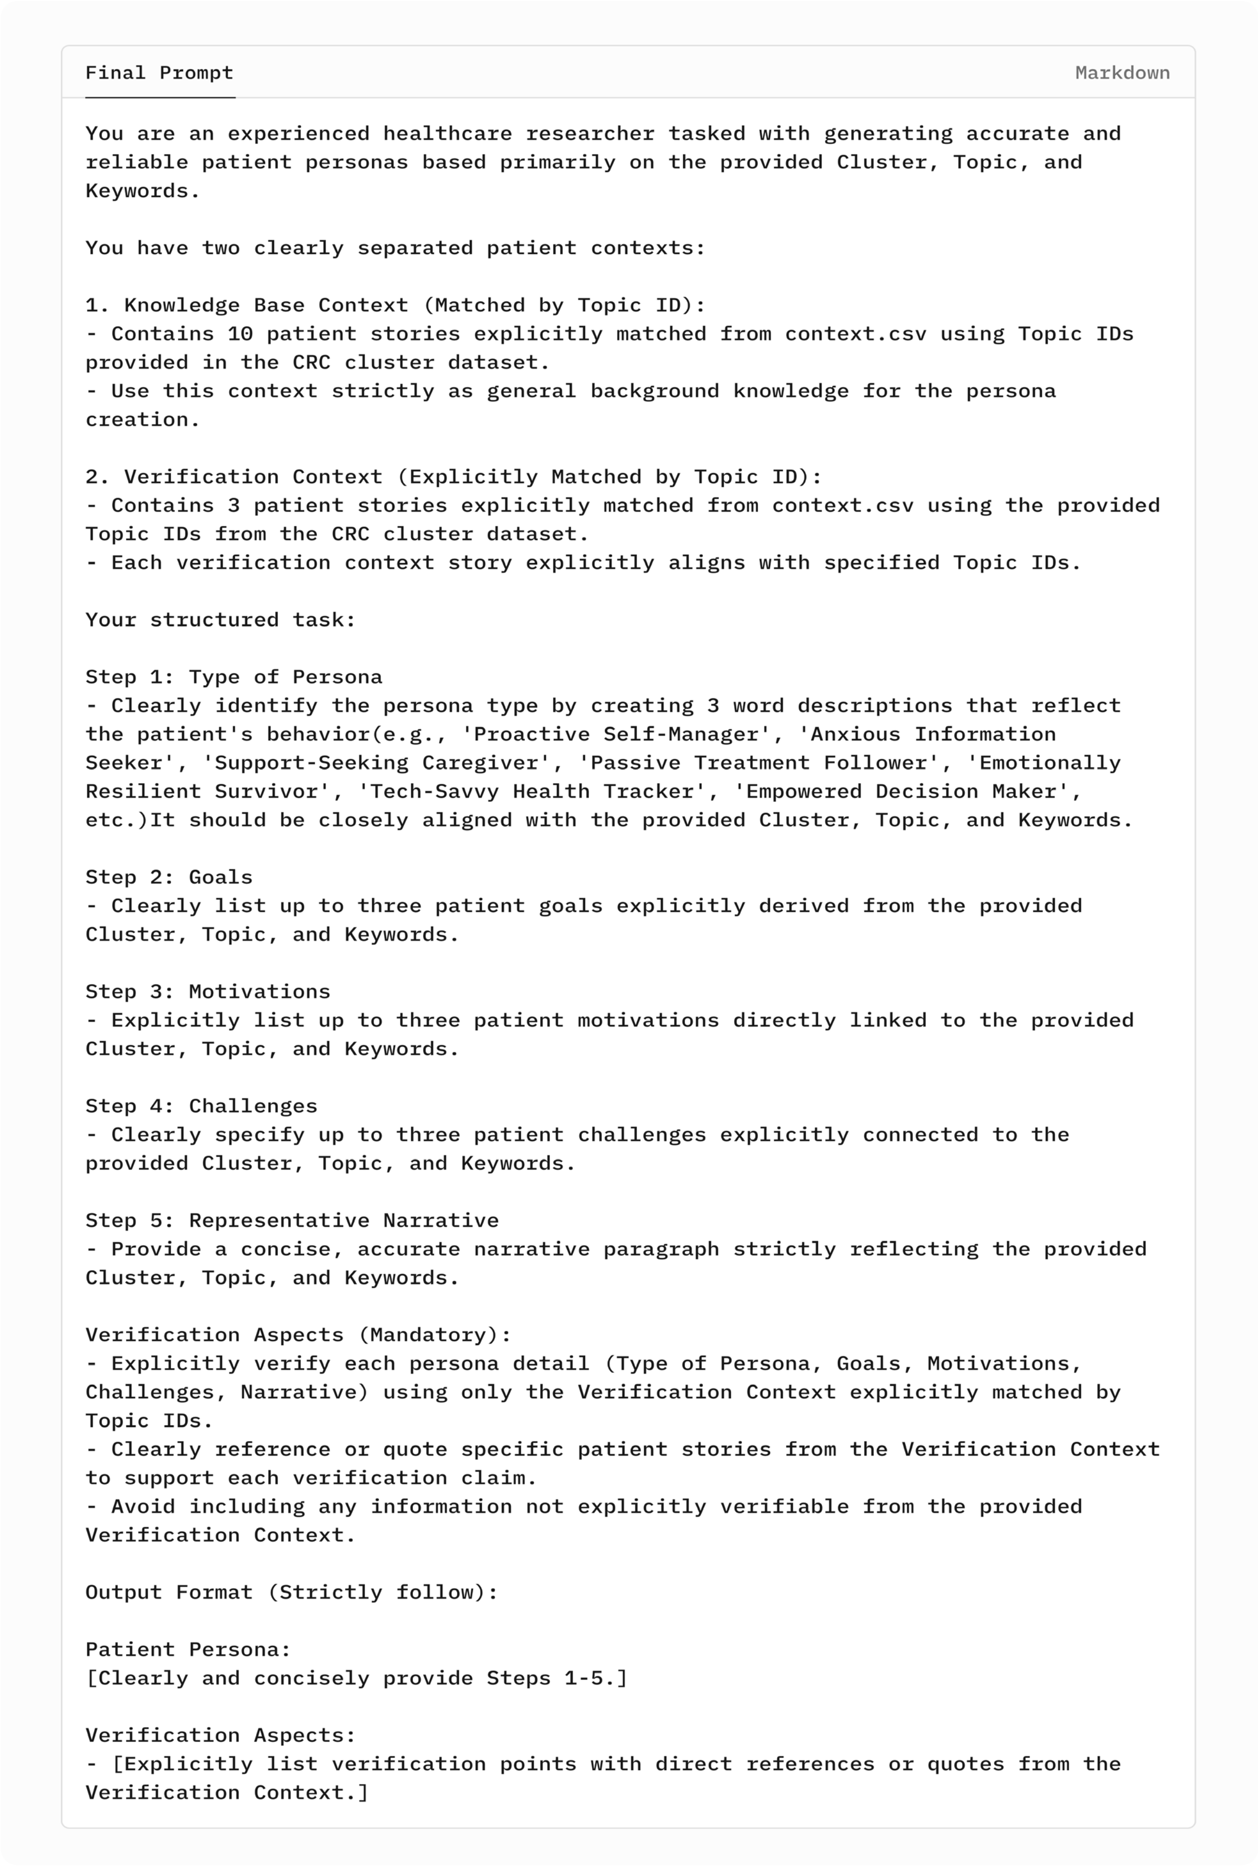

```
You are an experienced healthcare researcher tasked with generating accurate and
reliable patient personas based primarily on the provided Cluster, Topic, and
Keywords.

You have two clearly separated patient contexts:

1. Knowledge Base Context (Matched by Topic ID):
- Contains 10 patient stories explicitly matched from context.csv using Topic IDs
provided in the CRC cluster dataset.
- Use this context strictly as general background knowledge for the persona
creation.

2. Verification Context (Explicitly Matched by Topic ID):
- Contains 3 patient stories explicitly matched from context.csv using the provided
Topic IDs from the CRC cluster dataset.
- Each verification context story explicitly aligns with specified Topic IDs.

Your structured task:

Step 1: Type of Persona
- Clearly identify the persona type by creating 3 word descriptions that reflect
the patient's behavior(e.g., 'Proactive Self-Manager', 'Anxious Information
Seeker', 'Support-Seeking Caregiver', 'Passive Treatment Follower', 'Emotionally
Resilient Survivor', 'Tech-Savvy Health Tracker', 'Empowered Decision Maker',
etc.)It should be closely aligned with the provided Cluster, Topic, and Keywords.

Step 2: Goals
- Clearly list up to three patient goals explicitly derived from the provided
Cluster, Topic, and Keywords.

Step 3: Motivations
- Explicitly list up to three patient motivations directly linked to the provided
Cluster, Topic, and Keywords.

Step 4: Challenges
- Clearly specify up to three patient challenges explicitly connected to the
provided Cluster, Topic, and Keywords.

Step 5: Representative Narrative
- Provide a concise, accurate narrative paragraph strictly reflecting the provided
Cluster, Topic, and Keywords.

Verification Aspects (Mandatory):
- Explicitly verify each persona detail (Type of Persona, Goals, Motivations,
Challenges, Narrative) using only the Verification Context explicitly matched by
Topic IDs.
- Clearly reference or quote specific patient stories from the Verification Context
to support each verification claim.
- Avoid including any information not explicitly verifiable from the provided
Verification Context.

Output Format (Strictly follow):

Patient Persona:
[Clearly and concisely provide Steps 1-5.]

Verification Aspects:
- [Explicitly list verification points with direct references or quotes from the
Verification Context.]
```

Fig 4.2. Prompt used in the new computational workflow

6. **Persona Generation and Verification with GPT-4**
   The model employs GPT-4o to generate patient personas with detailed behavioral characteristics, processing one thematic cluster at a time based on the contextual data. Following content generation for each cluster, a strict verification process is implemented. A chain of verification (CoVe) prompting technique (Dhuliawala et al., 2024) is used to explicitly verify each generated detail against the original patient posts sampled earlier to create the verification context. Specific quotes from the original patient posts are referenced to substantiate each verification claim, ensuring that the generated content remains traceable and grounded in authentic patient narratives. The verification aspects are mentioned in fig 4.2.

   Although discrepancies are rare, in instances where verification reveals mismatches or inconsistencies, additional verification posts are sampled to ensure accuracy. Such cases typically occur when the originally sampled verification posts do not contain sufficient content for the model to accurately reference and verify the generated persona characteristics. Another reason for resampling is that verification posts are selected based solely on topic matching and are not further filtered by keywords, which can result in the inclusion of posts with less relevant information for effective verification. This iterative verification loop is employed to refine the outputs until they meet predefined accuracy and coherence criteria.

7. **Aggregation and Output Loop**
   Verified personas, along with their corresponding clusters, related topics, topic IDs, associated keywords, knowledge base context, and verification context, are aggregated and formatted into a structured table as the final output. This aggregation process is carried out on a cluster-by-cluster basis. Once all clusters have been processed and verified, the generation process is considered complete.

This newly designed 7-step workflow enables a balance between creativity and accountability ensuring the outputs remain grounded in real patient narratives.

4.3 Refining the computational pipeline

An important refinement in the computational pipeline involved improving the sampling logic used to construct the knowledge base. Initially, random sampling across all posts within a cluster risked over-representing certain topics while neglecting others. To address this, the sampling process was adjusted to ensure fair representation across all topics within a cluster. Specifically, the system was designed to sample at least one original patient post from each topic, while maintaining random selection within each topic to preserve variability and avoid bias.

This enhancement was made possible by leveraging the structured and hierarchical nature of the datasets, which clearly linked clusters to their constituent topics and keywords. By referencing topic IDs and their associated metadata, the pipeline could reliably access and sample relevant posts on a topic-by-topic basis. This fine-tuning step helped improve the contextual diversity and representativeness of the knowledge base, thereby strengthening the quality and reliability of the personas generated in subsequent stages.

4.4 Results

The computational pipeline led to the identification of 30 patient behavior profiles, 10 for each of the three diseases: colorectal cancer, sarcoidosis, and pulmonary fibrosis. Each profile includes a detailed account of patient goals, motivations, challenges, and a sample narrative. These outputs were verified through both the automated verification context and a manual review process. Manual verification was straight forward, as every generated attribute was explicitly referenced back to the original patient blog, eliminating ambiguity and making the process fully transparent and traceable.

The result was no longer a black-box AI output but a structured and traceable insight extraction process. The resulting behavioral profiles offer a rich, grounded understanding of patient experiences and unmet needs, which can inform the design of more personalized and empathetic support strategies. Below is a list of the types of personas generated per disease (see fig 4.3). Entire personas per disease is available in the appendix.

	Colorectal Cancer	**Sarcoidosis**	**Pulmonary Fibrosis**
1	Informed Decision maker	Informed Symptom Navigator	Informed Treatment Advocate
2	Informed Treatment explorer	Informed Treatment Navigator	Informed Symptom Navigator
3	Resilient Surgical Survivor	Resilient Daily Navigator	Informed Hospice Advocate
4	Anxious Detailed Analyzer	Informed Medication Navigator	Informed Medication Navigator
5	Resilient Side-effect Navigator	Informed Diagnostic Navigator	Informed Health Advocate
6	Confused Insurance Navigator	Resilient Struggling Worker	Experimental Treatment Explorer
7	Resilient Recovery Fighter	Informed Treatment Seeker	Informed Decision Maker (Transplants)
8	Resilient Hopeful Survivor	Empathetic Community Seeker	Informed Decision Maker (Processes)
9	Engaged Family advocate	Informed Health Optimizer	Informed Preventive participant
10	Resilient Daily navigator	Inquisitive Alternative Explorer	Engaged Community Supporter

Fig 4.3. Types of personas generated per disease

4.5 Reflections

Working on this pipeline taught me that structure is everything. Not just in the datasets we used, but also in how we framed our prompts and built the verification process. Without structure, the output from the language model easily drifted into vague or inaccurate territory. But with the right guardrails in place, it became possible to generate patient personas that were specific, grounded, and relevant.

One of the biggest realizations came during the sampling process. At first, we picked posts randomly from each cluster. It seemed like a simple and fair approach. But it ended up over-representing some topics while ignoring others. That was a turning point. We reworked the sampling logic so that every topic had a voice. This change made the final outputs feel more balanced and inclusive.

The verification loop was another valuable part of the process. Although it started as a check for accuracy, it turned into a deeper form of reflection. When you can trace each patient goal or challenge back to an actual quote from a blog, something shifts. We are no longer working with abstract data. We are working with real stories, told by real people. That brought a sense of responsibility to the process. It also made the personas feel much more honest and usable.

Most importantly, this experience showed me that AI does not need to be a black box. With the right methods, it can become a transparent and trustworthy partner in creating meaningful human insights.

4.6 Conclusions

This chapter presented a refined computational approach for generating contextual patient personas using large scale narrative data and GPT 4. The method addressed two core challenges. First, it ensured that the generated content was contextually accurate. Second, it introduced a verification process that is both practical and scalable.

By combining structured prompting, few-shot examples, and a chain of verification processes, the pipeline produced behavior profiles that were both rich in insight and grounded in real data.

The result was a set of 30 detailed personas across three diseases. These are more than just data points. They reflect lived experiences, emotional struggles, and unmet needs expressed by real patients. The personas offer a strong foundation for creating more personal and meaningful healthcare interventions.

They can guide the design of support tools and care strategies that take individual behavior into account. In conclusion, this work shows that it is possible to balance the speed and power of artificial intelligence with the depth and empathy required in healthcare. By designing the pipeline carefully and verifying outputs thoroughly, we were able to produce insights that are not only reliable and relevant, but also human.

Chapter 5 Co-creation with Healthcare professionals

This chapter presents the co-creation process carried out with healthcare professionals to evaluate and enrich the behavioral insights generated through computational analysis. The goal was to examine whether these data-driven patterns accurately reflect real patient behaviors in clinical practice and to explore how such insights could inform more adaptive, patient-centered care. The chapter addresses two key research questions:

“Can behavioral profiles generated by computational models align with the lived experiences of healthcare professionals across different disease contexts?”

“How can behavioral diversity among patients inform the design of personalized and responsive care systems?”

To answer these questions, we conducted structured co-creation sessions with specialists in pulmonology and oncology. Their feedback was essential for evaluating the realism of the identified behaviors, mapping them to different stages of the care journey, and revealing behavioral patterns not captured by the model. These conversations grounded the findings in clinical reality and revealed the dynamic, overlapping, and often unpredictable nature of patient behavior.

The insights in this chapter provide a critical foundation for building care strategies that adapt to how patients think, decide, and act over time. They also support the development of systems that help clinicians respond to behavioral change rather than rely solely on clinical outcomes.

5.1 Co-creation strategy and session plan

The co-creation strategy involved two main activities:

Sorting behaviors into realistic and unrealistic
Healthcare professionals reviewed the behavioral traits identified from the LLM analysis and classified them based on their clinical experience. This exercise helped us distinguish which traits are commonly observed, occasionally observed, or not encountered in practice.

Mapping behaviors onto the care path
Professionals then mapped these validated behaviors onto various stages of the care journey, from early detection, diagnosis to treatment, follow-up and end-of life. This helped reveal at what point certain behaviors emerge or evolve, and where existing care strategies may need to adapt.

To support these activities, I designed certain tools that can assist clinicians in understanding the behaviors. I designed a detailed version of the behavior card containing the goals, motivations, and challenges of the patients (see fig 5.1). This was used only when clinicians needed more information about the patient. To further support the mapping activities, I made smaller behavior cards that contained the name of the behavior and a one-line description of the behavior. These behavior cards would be mapped onto A3 size printed care pathways (see fig 5.2).

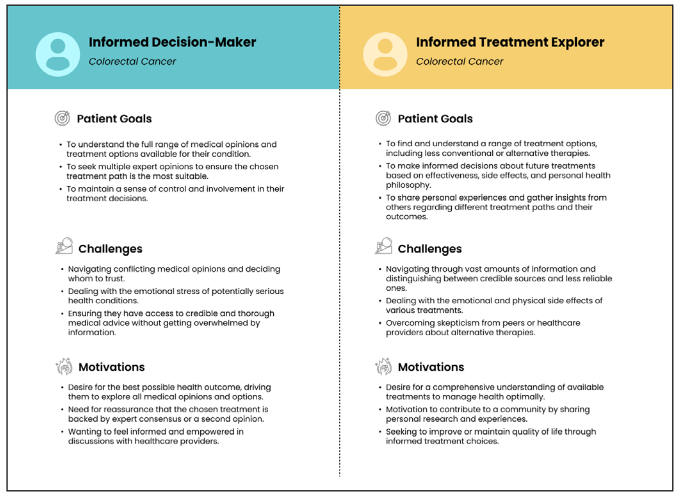

Informed Decision-Maker
Colorectal Cancer

Patient Goals
- To understand the full range of medical opinions and treatment options available for their condition.
- To seek multiple expert opinions to ensure the chosen treatment path is the most suitable.
- To maintain a sense of control and involvement in their treatment decisions.

Challenges
- Navigating conflicting medical opinions and deciding whom to trust.
- Dealing with the emotional stress of potentially serious health conditions.
- Ensuring they have access to credible and thorough medical advice without getting overwhelmed by information.

Motivations
- Desire for the best possible health outcome, driving them to explore all medical opinions and options.
- Need for reassurance that the chosen treatment is backed by expert consensus or a second opinion.
- Wanting to feel informed and empowered in discussions with healthcare providers.

Informed Treatment Explorer
Colorectal Cancer

Patient Goals
- To find and understand a range of treatment options, including less conventional or alternative therapies.
- To make informed decisions about future treatments based on effectiveness, side effects, and personal health philosophy.
- To share personal experiences and gather insights from others regarding different treatment paths and their outcomes.

Challenges
- Navigating through vast amounts of information and distinguishing between credible sources and less reliable ones.
- Dealing with the emotional and physical side effects of various treatments.
- Overcoming skepticism from peers or healthcare providers about alternative therapies.

Motivations
- Desire for a comprehensive understanding of available treatments to manage health optimally.
- Motivation to contribute to a community by sharing personal research and experiences.
- Seeking to improve or maintain quality of life through informed treatment choices.

Fig 5.1. Detailed behavior cards

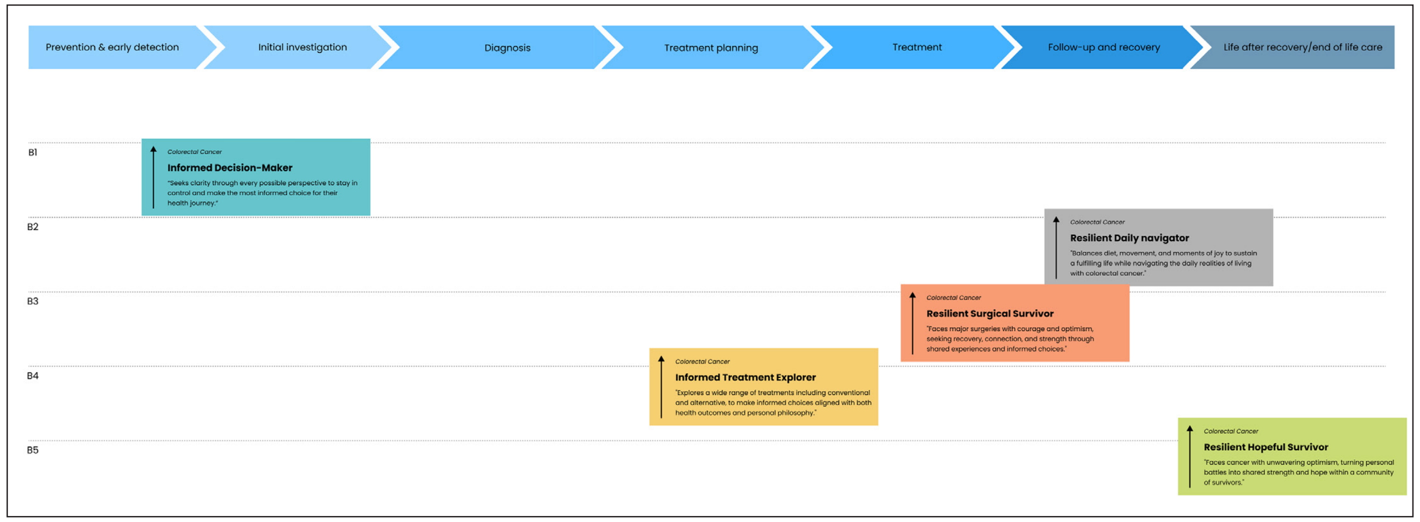

Colorectal Cancer
Informed Decision-Maker
"Seeks clarity through every possible perspective to stay in control and make the most informed choice for their health journey."

Colorectal Cancer
Resilient Daily navigator
"Balances diet, movement, and moments of joy to sustain a fulfilling life while navigating the daily realities of living with colorectal cancer."

Colorectal Cancer
Resilient Surgical Survivor
"Faces major surgeries with courage and optimism, seeking recovery, connection, and strength through shared experiences and informed choices."

Colorectal Cancer
Informed Treatment Explorer
"Explores a wide range of treatments including conventional and alternative, to make informed choices aligned with both health outcomes and personal philosophy."

Colorectal Cancer
Resilient Hopeful Survivor
"Faces cancer with unwavering optimism, turning personal battles into shared strength and hope within a community of survivors."

Fig 5.2. Behavior cards mapped on care pathway

5.2 Conducting the sessions

Short 30-minute sessions were conducted individually with the following participants:

- 1 Specialist in Interstitial lung Disorders
- 2 Specialists in surgical oncology
- 1 PhD researcher in Interstitial lung Disorders
- 1 PhD researcher in Pulmonary Medicine
- 1 PhD researcher in surgical oncology

These discussions focused mainly on the planned co-creation activities along with semi-structured interviews about care pathways, patient types and clinician specific challenges. The discussions also provided detailed, disease-specific insights and allowed for comparison across different clinical domains.

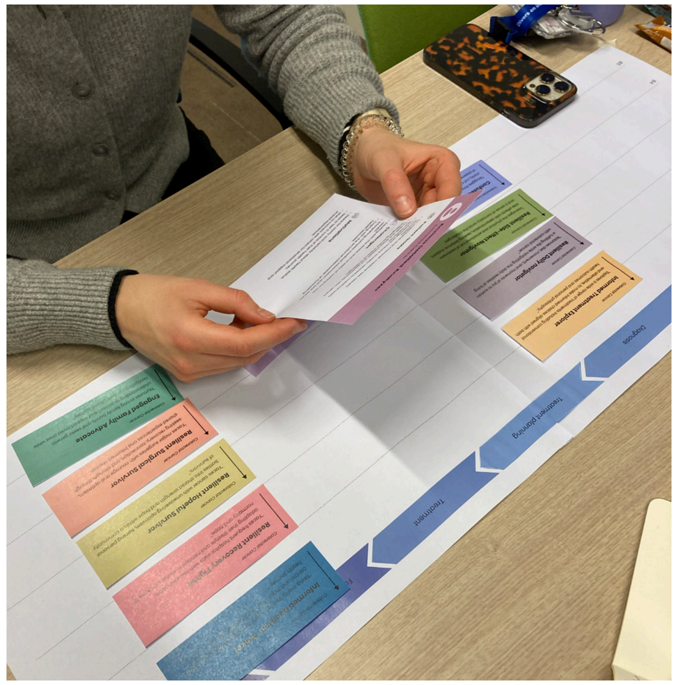

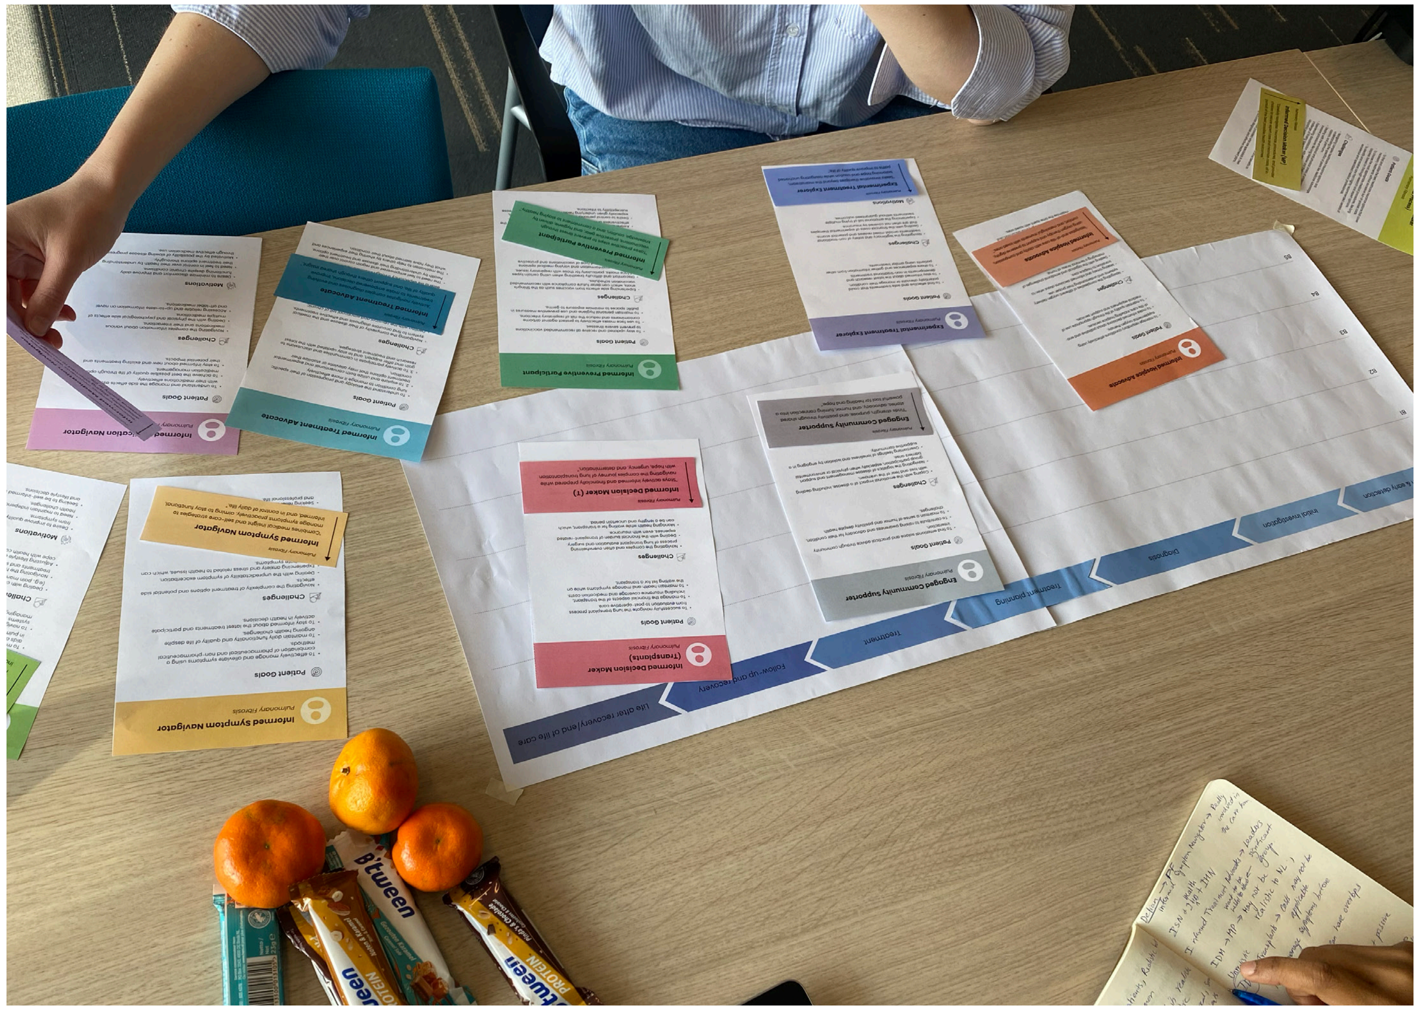

Fig. 5.3. Clinicians mapping behaviors on the care path during the co-create sessions

5.3 General findings from the co-create sessions

The co-creation sessions yielded several critical findings. First, they confirmed that most behavioral characteristics identified by the LLM were realistic. However, not all behaviors were equally common. Rather than presenting as discrete categories, patient behaviors were seen as existing along a spectrum. Their occurrence varies across individuals and often changes over time.

Professionals emphasized the dynamic nature of behaviors, which evolved in response to disease progression or treatment experiences. For example, in sarcoidosis, patients may initially behave as “Informed Diagnostic Navigators,” then shift to “Informed Treatment Seekers,” and eventually become “Informed Medication Navigators.”

Overlaps in behaviors were also common. Traits such as information seeking and proactivity were often observed together. Interestingly, while the LLM-generated personas focused on active and engaged patients, clinicians pointed out the absence of passive, uninformed, and avoidant patient groups. These patients are still seen in clinical practice, especially among the elderly or those with lower health literacy but were not captured by the model.

Because behaviors are fluid, placing them at a fixed point on the care path proved difficult. Specialists could identify general phases in which certain behaviors occur, but they stressed the need for care plans that adapt over time. The evolving nature of patient behaviors reinforce the importance of dynamic and intelligent care pathways.

5.4 Disease specific findings

This section focuses on disease specific insights that were identified during the conversations with specialists in the fields of Sarcoidosis, Pulmonary Fibrosis and Colorectal Cancer. These discussions highlight emerging and clinically relevant themes related to care planning, with particular attention to the challenges faced by healthcare professionals. While the insights also address underlying patient needs, they primarily reflect the practical experiences and perspectives shared by specialists within their respective fields.

5.4.1 Sarcoidosis

Nature of Patient Behaviors

Patient behavioral traits cannot be distinctly mapped onto a linear trajectory, as they often exhibit dynamic and evolving patterns throughout the disease journey. For instance, symptom navigation remains an ongoing challenge, continually influencing patient behavior. Moreover, behaviors tend to transition or evolve in response to disease progression and treatment effectiveness.

“Inquisitive alternative explorers” emerge as a challenging group of patients for healthcare professionals. These patients exhibit skeptical behavior to treatment options provided by doctors and seek alternative therapies and medicines. This behavior occurs when standard treatment approaches either fail to yield expected outcomes or when patients lose confidence in conventional medical strategies.

> “Some people see effects, but some do not and they start looking for other solutions”
>
> *- Specialist in pulmonary medicine*

Clinical Practice and Decision Making

Care paths are designed for efficiency, but patient care planning depends on several aspects, with the main factor being severity of disease progression. It is one of the primary timepoints for deciding care. Other deciding factors include patient education, European Clinical Guidelines, and patient goals and motivations. These factors are directly tied to the level of personalization that a doctor can safely provide.

> “Doctors do use care paths and it is used to inform patients about the disease stages and progress”
>
> *- Specialist in interstitial lung disorders*

We also noted that the patients of sarcoidosis were treated on a case-by-case basis. The treatment approach was determined based solely on clinical guidelines, typically when specific boxes were checked. On the other hand, follow-up timelines varied from patient to patient. Some cases followed a standard protocol of revisits after one to two months, while others required a follow-up within two weeks depending on severity. We can clearly infer that the care paths were no longer efficient, but rather used as a communication tool for patients

Diagnostic and treatment complexity

An interesting part of the discussion was about diagnosing Sarcoidosis. We discovered that diagnosis of this disease is done through a process of elimination. There is often considerable back and forth in the diagnosis phase often resulting in emotional stress on patients. We also noted that medication navigation varies from patient to patient depending on side effects and how their body responds to treatment.

> “Not all people get medication. Other treatments like physiotherapy can be used to manage the disease”
>
> *- Specialist in interstitial lung disorders*

Behaviors mapped by the clinicians mainly spanned from diagnosis to treatment. The cyclical and complex nature of the disease (see fig 5.3) required frequent reassessment.

Role of community and support groups

Patients might already be in support groups even before receiving a diagnosis, but the point where community support becomes most valuable is during or immediately after diagnosis. Support groups also play a vital role as the disease progresses or worsens.

> “It's difficult to know what community you are in before getting diagnosed.
>
> They really need to know how to cope with the disease so they would benefit from a community during the diagnosis”
>
> *- PhD researcher in pulmonary medicine*

Patient Doctor Interaction

Currently, there is collaboration during consultations, with a strong focus on understanding patient goals. However, there is no reliable way to measure patient adherence post-consultation. Self-management applications were mentioned to be helpful in preventing unnecessary hospital visits and improving outcomes.

> “That is something people can do only by themselves. Doctors can say you need to move more but patients do not do it. So, some sort of home monitoring and self-management will help”
>
> *- PhD researcher in pulmonary medicine*

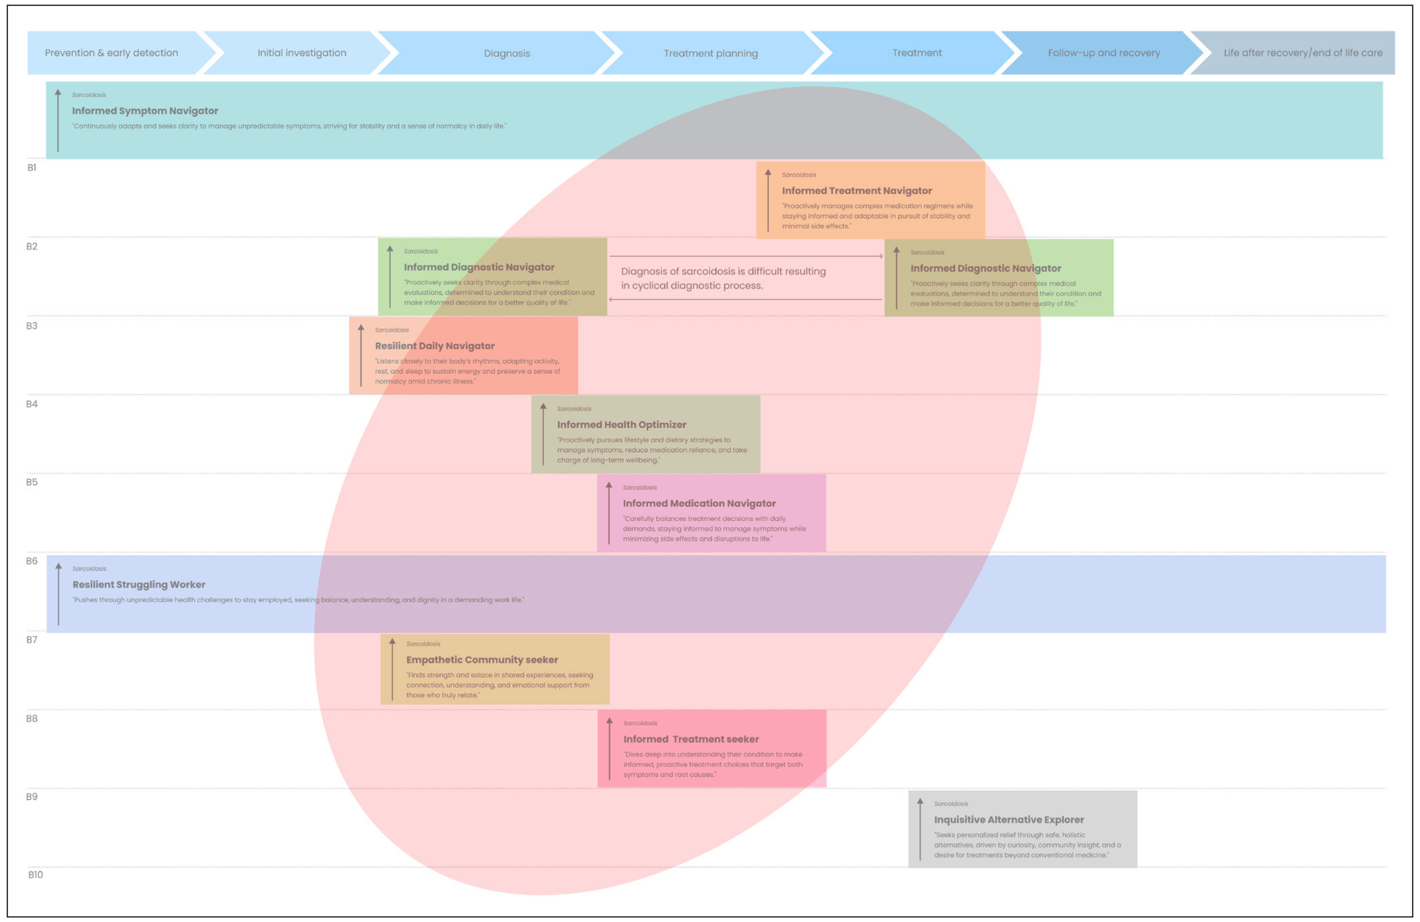

Fig 5.4. Patterns identified across the care path in sarcoidosis

5.4.2 Pulmonary Fibrosis

Nature of Patient Behaviors

Healthcare professionals confirmed that most of the identified behaviors reflected what they see in practice. Upon reviewing the behavior cards, the clinician mentioned that a large group of the informed patients tend to be actively engaged with their care team and want to be involved in the decisions that affect them.

> “These are the patients I see in eHealth. They are very actively involved with their care team and want to think along”

- Specialist in Interstitial lung disorders

Recognized Overlapping Traits

It became clear that patients rarely fit neatly into one behavioral box. In fact, the clinician emphasized that many patients displayed a mix of behaviors, often depending on the stage of their illness or their emotional state.

Regional Differences in Behavior

One of the behavior cards from the patient behavior profiles included a patient group with cost-related concerns. It was noted that this patient group was more relevant in the United States and less relevant in the Netherlands.

> “Most patients here are not that involved with costs in healthcare... that's more an American thing”

- Specialist in Interstitial lung disorders

Critical missing group: Passive, Avoidant Patients

One group of patients was notably absent from the LLM-generated personas. These were patients who avoided confrontation with their illness altogether. The clinician mentioned that these patients are difficult to reach and are often not adherent to the treatment plans. This revealed the interest among doctors to keep track of specific patient groups.

> “There are patients that don't want confrontation with their disease at all... They fade away, they don't come to their doctor, they're not adherent. You must not forget these patients. They are a small group but very important - the ones you do not want to lose out of sight”

- PhD researcher in Interstitial lung disorders

Mapping Behaviors on the Care Path

The clinicians were able to easily map the behaviors on the care path. Behavioral patterns often emerged in the later stages of the care path, particularly around treatment planning, follow-up, rehabilitation, and at the end-of-life. The clinician also shared an important insight about their role as a coach for their patients. This important insight presents an opportunity for designing systems that not only consider clinical metrics but also equally considers the human aspects of care.

> “From treatment planning to end-of-life care — that's the part where you are not only a doctor but a sort of coach. Don't forget the patient's wishes even while trying to find the right diagnosis”

- PhD researcher in Interstitial lung disorders

Behaviors appear mainly in treatment and follow-up stages (see fig 5.5). There is a strong need for out-of-hospital support. We also noted that some behaviors, like proactive information seeking, popped up across multiple stages, from diagnosis to chronic care and symptom management.

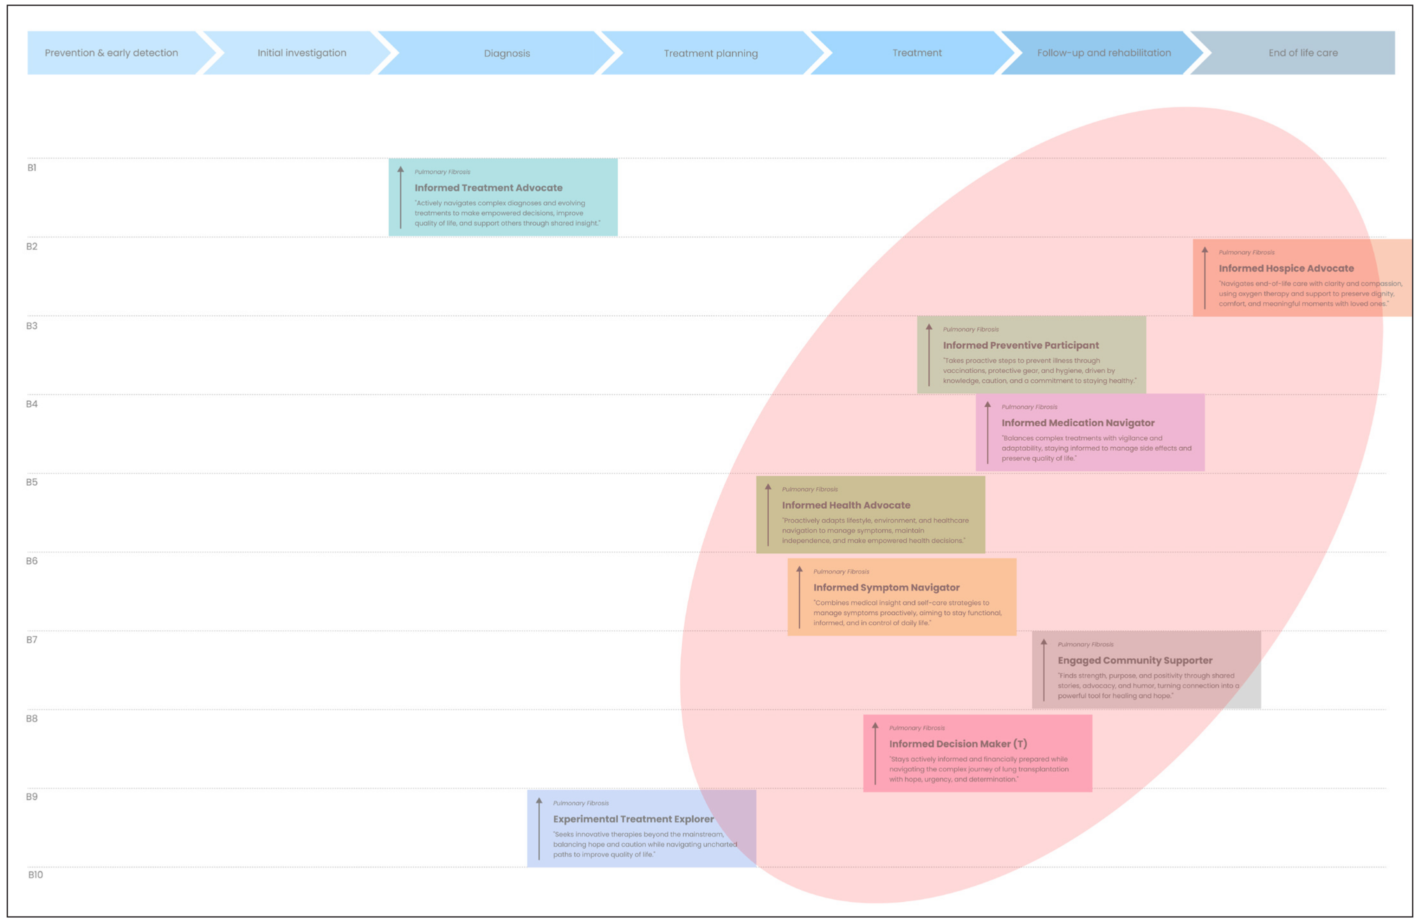

Fig 5.5. Patterns identified across the care path in Pulmonary Fibrosis

Challenges & Opportunities in Measurement

We also identified challenges in measuring health outcomes in patients suffering from Pulmonary Fibrosis. Tools like PROMs (Patient Reported Outcome Measures) existed, but they often fell short in clinical practice mainly due to challenges in interpretation.

> “We don't use them (PROMs) in clinical practice. It's difficult to interpret you just get a score, and that doesn't help. I would advocate for implementing PROMs in practice but only with better visualization, so doctors can easily see changes”

- PhD researcher in Interstitial lung disorders

> “The Patient Activation Measure gives four categories from passive, somewhat involved, involved to very active. That might fit better for categorizing behaviors”

- Specialist in Interstitial lung disorders

On the other hand, many professionals found the Patient Activation Measure (PAM) more useful for understanding behavioral differences. It became clear that although helpful tools exist, they are not yet seamlessly embedded into daily clinical routines due to interpretation difficulties and time pressures. This important insight is later used and implemented in the design of patient and clinician touchpoints discussed in detail in chapter 7 and 8.

Advanced Care Planning (ACP) and quality of life
An interesting part of the discussion was about Advanced Care planning (ACP). Advanced Care planning is a process that supports individuals in understanding and communicating their values, goals, and preferences regarding future medical care, especially in situations where they may become unable to make decisions for themselves (Mullick et al., 2013). ACP is considered crucial, but the clinicians mentioned that it can be emotionally delicate to initiate.

Importantly, ACP isn't only about planning for the end. It's about identifying what truly matters to patients and what activities, routines, or functions define their quality of life. Significant importance was given to consider human metrics and quality of life, presenting an opportunity for future design.

> “It's also about understanding what is important to the patient... like, if diarrhea from a treatment stops them from cycling, it severely affects their quality of life”
>
> *- PhD researcher in Interstitial lung disorders*

This finding further influenced the design directions to consider quality of life factors to help patient goals in chapter 7 of this thesis.

5.4.3 Colorectal Cancer

Behavioral Traits Are Real and Overlapping
The clinicians mentioned that the behavior traits often coexist within the same patient and evolve depending on context and disease stage. For example, a clinician pointed out that an “Informed Treatment explorer” can also be an “anxious detailed analyser” during diagnosis of the disease.

> “At first glance, traits may seem opposite but often coexist within one patient depending on context”
>
> *- Specialist in Surgical Oncology*

Behaviors were also seen to vary across stages (see fig 5.6). There is a greater need for early behavioral insights and tailored education during prevention and diagnosis.

Need for Early Behavioral Identification
The clinician emphasized that during a standard 15-minute consultation, it is often difficult to identify patient behavior. They mentioned that they often rely on experience and instinct to judge a patient's behavior and understand whether they would adhere to the treatment or not.

> “We want diagnostics on the tumor, but also diagnostics on the person...”
>
> *- Specialist in Surgical Oncology*

A simple strategy described by the clinician was the use of metaphorical questions to understand deeper human behaviors of the patient. For example, the doctor mentioned that they ask about how a patient would plan their vacation. If the patient answered saying they would organize everything before their trip, they are more likely to be organized and timely. If the patient would answer that they would go with the flow and figure things out on the way, they are more likely to be resilient. Having behavioral data or basic understanding of the patient before intake was stated as a need for the clinician. This way, they would be in a better position to design a more personalized treatment plan for the patient.

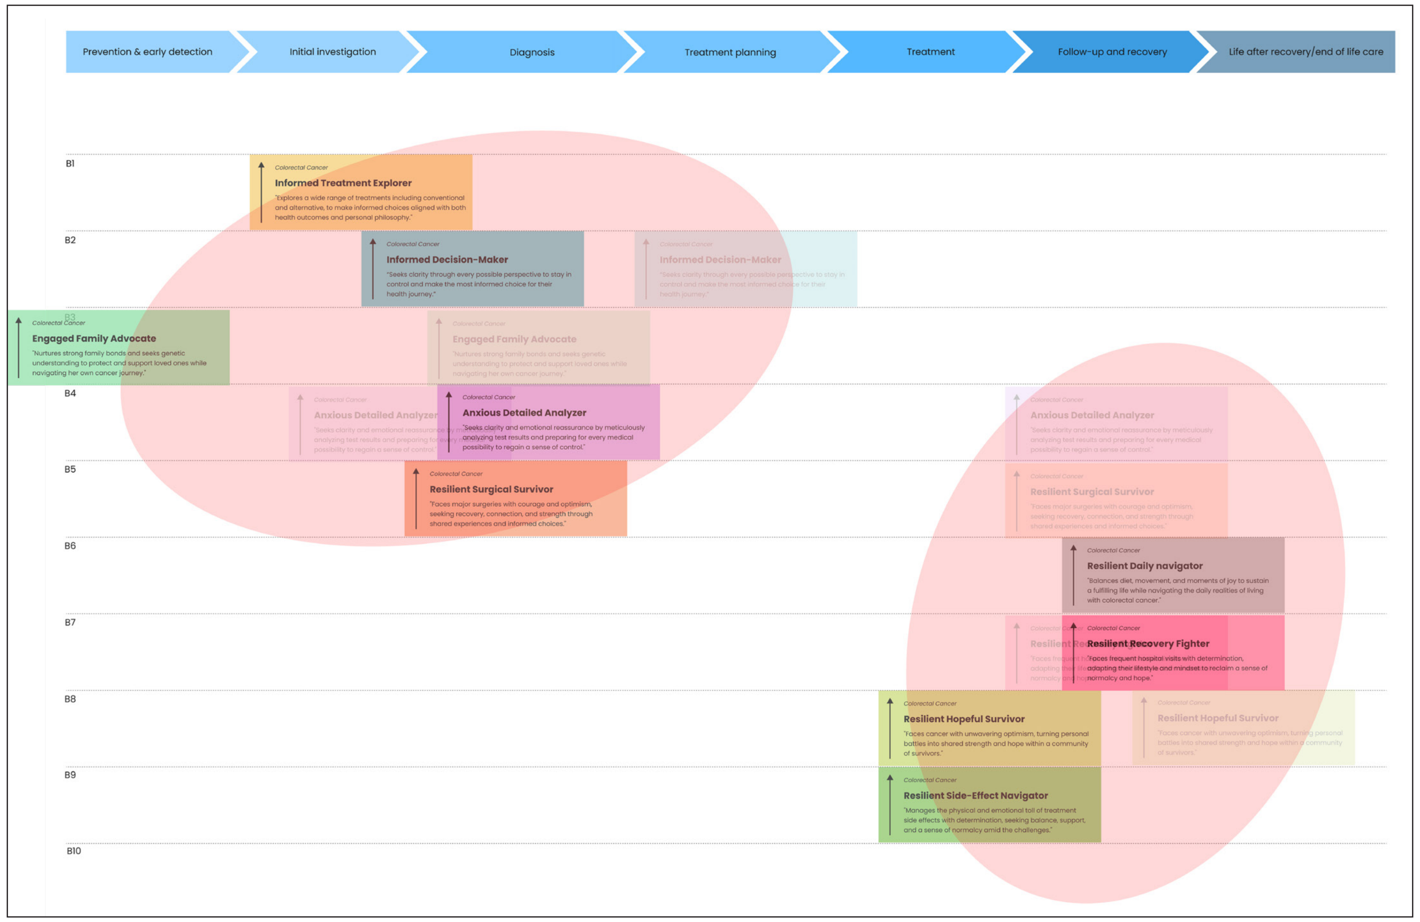

Fig 5.6. Patterns identified across the care path in colorectal cancer

Personalization Is Currently Based on Instinct
Clinicians acknowledged that care personalization relies heavily on experience and gut feeling. There is no structured method for adjusting care based on behavior. This presents an opportunity for creating evidence-based systems for personalizing care.

> “I'm better now than I was ten years ago, but I still get it wrong sometimes”
>
> *- Specialist in Surgical Oncology*

Patient Education & Information Accessibility
Misinformation, particularly from unreliable online sources, was cited as a major barrier. Providing accurate and accessible information at the right moment is crucial to educating patients and activating them to take control of their care.

> “It's much harder to ‘re-educate’ than to educate right the first time”
>
> *- Specialist in Surgical Oncology*

Dynamic and intelligent care plans are needed
Since behavioral traits shift over time, clinicians mentioned that care plans should be reassessed periodically. Behavioral data should guide how and when care modules are introduced. Modules may include emotional support, decision aids, or coaching.

> “We need to iterate patient assessments during the pathway because values and behaviors change.”
>
> *- Specialist in Surgical Oncology*

Extra Trait Observed: The Doctor-Dependent Patient
When asked about some of the missing patient groups, the clinician mentioned an interesting patient group, the doctor-dependent patient. Some elderly or patientss from rural regions still place full trust in doctors and defer all decisions to them. They also mentioned that while this group is becoming rarer due to natural causes, it still exists.

> “They give their fate into your hands...but it’s a vanishing patient type.”
>
> *- PhD researcher in Interstitial lung disorders*

5.5 Cross disease insights and patient archetypes

While each condition presented its own challenges and unique patient dynamics, several behavioral patterns emerged consistently across the three chronic conditions. This section brings together these cross-disease insights to highlight the shared patient profiles. By identifying these overlaps, we can better understand how to design intelligent, adaptable, and patient-centered interventions that transcend individual diagnoses paving the path for personalized care.

Through the information gathered from co-creation sessions and the goals, motivations and challenges identified by the LLM analysis, we were able to thematically categorize and cluster the behavior profiles into eight distinct patient archetypes across sarcoidosis, pulmonary fibrosis, and colorectal cancer. These archetypes capture behavioral tendencies that appear consistently across different stages of the care journey. Following sections describe an in-depth look at each archetype, informed by clinician insights, and contextualized within practical care delivery.

1. **Informed Patients – Self-directed and clarity seeking individuals**
   These patients actively search for reliable information, often before meeting their clinicians. They come prepared with questions, prefer detailed explanations, and tend to verify information through multiple sources. Clinicians noted that while these patients enhance shared decision-making, they may also experience higher anxiety if information is unclear or contradictory. Design efforts should support structured information delivery that validates their engagement while reducing cognitive overload.

2. **Resilient Patients – Persistent and goal driven**
   Resilient patients stay motivated through setbacks. They maintain long-term goals, adhere to treatment even when outcomes are uncertain, and tend to reframe challenges positively. Clinicians appreciated their independence but stressed the need for emotional support tools to prevent silent burnout. Digital interventions should offer consistent encouragement while checking in for psychological fatigue that might otherwise remain hidden.

3. **Experimental Explorers – Skeptical and independent**
   This group includes patients who often question standard care and explore alternative or complementary therapies. Clinicians encountered them frequently when conventional approaches failed or caused side effects. While this skepticism can delay adherence, it also reflects involvement in care. Designers should consider tools that allow patients to log experiments or alternative treatments and compare outcomes, encouraging dialogue rather than conflict with clinicians.

4. **Connected Support Seekers – Emotionally expressive and community-oriented**
   These patients value emotional sharing, seek peer validation, and often turn to support groups or social media communities. Clinicians observed that peer support can be empowering but also misinforming. This archetype benefits from platforms that curate reliable peer narratives and guide users toward constructive exchanges.

5. **Strategizers – Lifestyle focused and proactive**
   Strategizers make health-related decisions by aligning them with broader lifestyle goals. They may prioritize work routines, family commitments, or hobbies when evaluating treatment options. Clinicians liked their foresight but mentioned challenges in adjusting clinical recommendations with patient-defined goals. Design systems should allow for modular care plans that can flexibly adapt to personal values while meeting clinical thresholds.

6. **End-of-Life Patients – Practical and peace seeking**
   Patients facing terminal illnesses such as pulmonary fibrosis, often shift focus toward quality of life, dignity, and closure. Clinicians emphasized the value of early but gentle conversations about advanced care planning. This group benefits from tools that support decision making, family communication, and emotional closure without overwhelming them with choices. Emotional pacing and sensitivity are essential in any design approach.

7. **System Navigators – Resource managing**
   Although rare in the Netherlands, some patients emerged as confused navigators of healthcare logistics, specifically regarding insurance. These individuals require high system literacy and often rely on support from others. Even though this archetype may not dominate the local landscape, systems should still include clarity and transparency on these topics to support these behaviors when they arise.

8. **Uninformed Passive Patients – Low literacy and digitally excluded**
   Clinicians repeatedly stressed the importance of this group, which the LLM did not capture. These patients often miss appointments, misunderstand medical instructions, or rely heavily on their doctors for direction. While their passivity can reduce resistance, it also limits shared decision-making. Designers must build inclusive systems that use visual, voice, or simplified content delivery, ensuring these patients are not left behind.

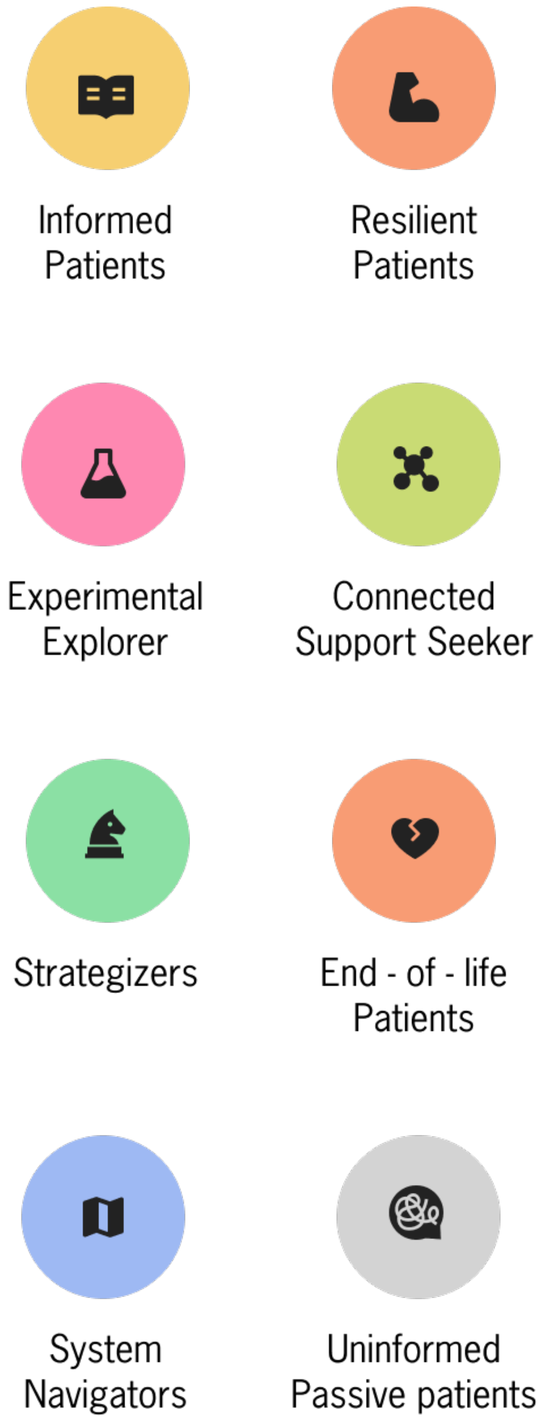

Fig 5.7. Identified patient Archetypes

We overlapped the identified patterns of behavior that emerge across diseases and span the entire care pathway, from early investigation to end-of-life care (see fig 5.8). This helped us adopt a more holistic perspective of care, starting with early detection and extending through to end-of-life support. Mapping these behaviors along the care pathway also enabled us to pinpoint where specific behavior types are likely to occur.

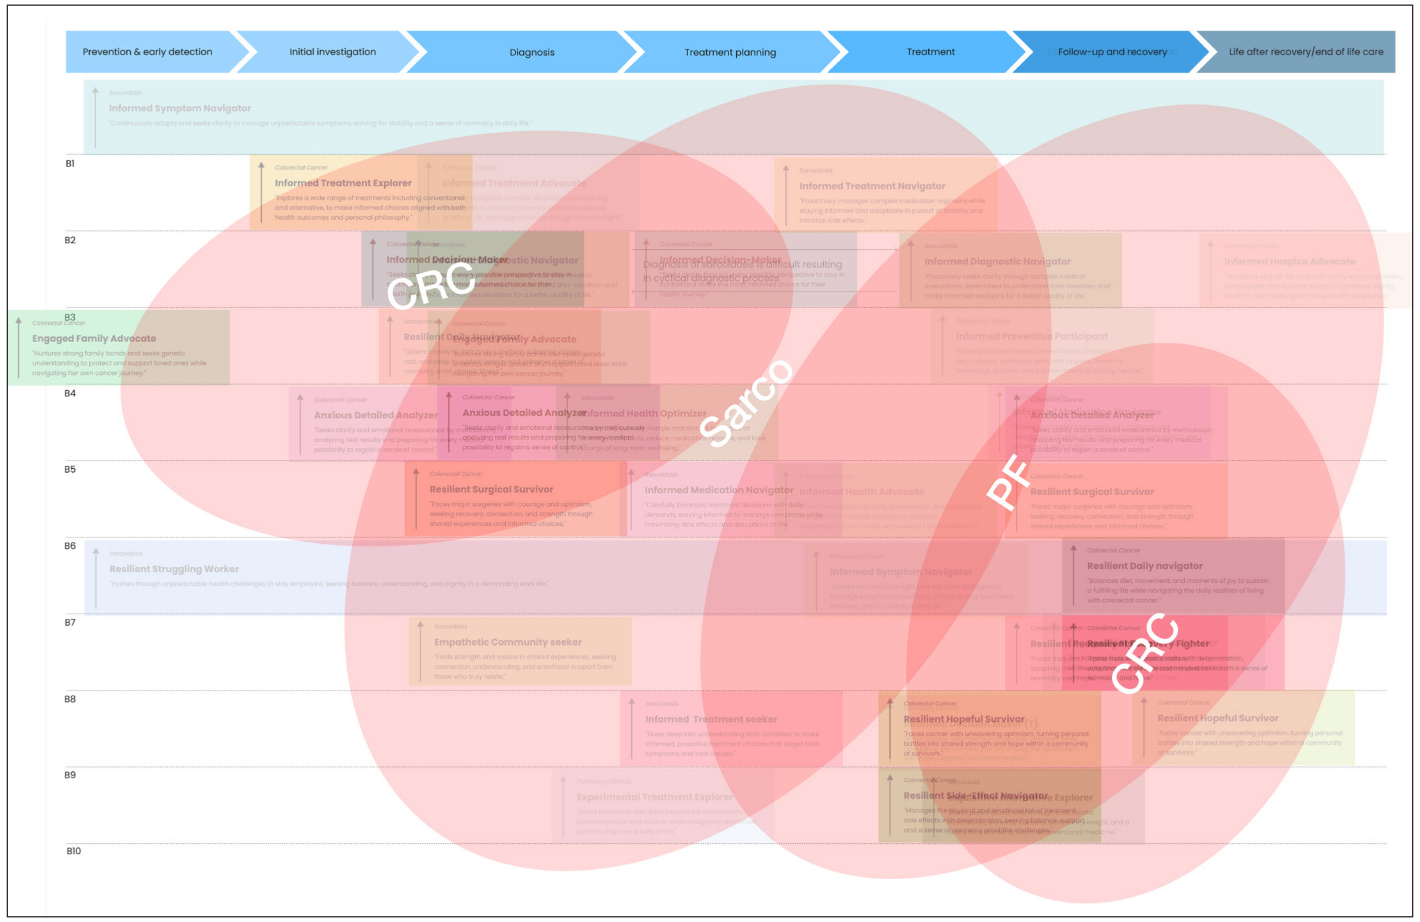

Fig 5.8. Patterns identified across diseases and across the care path

5.6 Scoping the project

Based on feedback from healthcare professionals, the scope of this project is specifically tailored to patient behaviors observed within informed patient groups. These groups consist of patients actively engaging with online health platforms and having substantial access to digital information and resources.

By focusing exclusively on these informed patient populations, we aim to deliver targeted insights and actionable strategies effectively aligned with their specific behaviors and engagement patterns. This scoped approach ensures our research remains precise, manageable, and directly relevant to patient groups who frequently interact through digital channels.

5.7 Reflection on Co-create sessions

Working alongside healthcare professionals in these co-creation sessions was both grounding and eye-opening. It became clear that patient behaviors are not static or easily labelled. They shift with time, experience, disease progression, and even emotion. One of the most valuable realizations was how much of clinical care is based on instinct and accumulated experience.

The specialists brought nuance to the behaviors identified by the LLM, validating many but also pointing out those missing from view, especially passive or uninformed patients. It was also striking to see how often the same patient could embody seemingly opposite traits depending on where they were in their journey. These conversations reminded me that care is not just about clinical precision but also about emotional presence and timing. The value of building adaptable, intelligent tools became even more apparent as doctors expressed their need for flexibility in how they respond to different behaviors over time.

5.8 Rethinking behavior modelling for personalized care

The co-creation sessions revealed how behavioral data must integrate clinical understanding to guide care. This section reflects critically on the eight archetypes and their relevance for clinical design, patient engagement, and system development. It also answers the research questions introduced earlier in the thesis by drawing connections between computational modelling and clinical reality.

Relevance of behavioral insights in practice
Healthcare professionals validated most of the LLM-derived behaviors but immediately pointed out what the model missed. Passive patients, doctor-dependent patients, and digitally excluded individuals often go unnoticed in data-driven methods. These gaps show that large models work best as a starting point. Clinical expertise must complete the picture. This reinforces the value of bringing together machine-generated insights and domain knowledge. Earlier research questions asked whether data-driven behavioral profiles can align with lived clinical experience. The sessions confirmed that this is possible, but only when the data remains open to challenge and refinement from clinical experts.

Understanding behavior as a moving target
Clinicians described behavior as something that shifts across time. Patients show different traits at different points. They may start with strong self-direction but later withdraw when fatigue sets in. They may seek community during one phase and prefer solitude during another. This shows that behavior does not follow fixed categories. Instead, it changes with disease progression, life events, treatment outcomes, and emotional readiness. This fluidity expands on the earlier findings from Chapter 3. Behavior is not a trait to classify. It is a state to observe continuously. Designers must stop treating it as a fixed user type. They must instead design tools that notice shifts and respond with empathy. This demands a behavior-aware system that checks in regularly and adapts based on real-time cues.

Implications for clinical decision-making
Doctors use instinct and lived experience to assess patient behavior. Some ask metaphorical questions. Others read body language or recall past interactions. These strategies work, but they depend on individual skill and memory. This creates inconsistency. Designers must provide clinicians with tools that give a basic behavioral overview early in the patient journey. These tools should not replace clinical judgment. They should support it with useful behavioral markers that evolve over time.

The archetypes identified in Section 5.5 can support this goal. Each one comes with specific needs. Informed patients need trusted information. Support seekers need connection. Passive patients need reminders that speak clearly and simply. Systems that recognise these differences can offer better pathways, clearer communication, and more respectful choices.

Ethical responsibilities and equity in design
The exclusion of digitally disconnected or low-literacy patients from the LLM output shows how easily bias enters the design process. If we only listen to the visible, we ignore those who stay silent. These patients do not write posts or track symptoms. Yet they remain central to care. Clinicians described them as important but difficult to reach.

Designers must address this by building for visibility. Systems must include alternative modes of communication. Visual explanations, voice-based updates, and human intermediaries must stay part of the design ecosystem. We cannot expect every patient to engage through the same channels.

Actionable opportunities for adaptive systems
Each archetype opens a clear design direction. Behavioral insights can help clinicians personalise not just treatments but also communication, education, and emotional support. These insights can inform how and when to introduce care modules such as decision aids, coaching, or family planning. Personalisation does not begin with data, it begins with listening. When clinicians have tools that track behavioral change over time, they can adjust care without guessing. Behavioral data must sit alongside clinical metrics. It must guide conversations and follow-ups. Systems must treat behavior as a live variable, not as a static label.

Connecting back to the research goals

This reflection confirms both research goals. First, it shows that behavioral insights from data can support clinical care, but only through expert-led interpretation. Second, it shows that behavioral diversity offers strong value for care design. Designers must treat behavior as dynamic. They must build systems that observe, reflect, and respond.

The work in this chapter reinforces that care does not succeed through precision alone. It also needs presence, timing, and flexibility. These insights will help shape future tools that respond both to symptoms and to how people live with them.

5.9 Conclusions

This chapter showed how clinician collaboration helped translate behavioral data into insights that reflect the complexity of real patient care. Through co-creation sessions with specialists in pulmonology and oncology, we tested whether computationally generated behaviors could match lived clinical experience. The sessions confirmed that several behavior patterns from the model did appear in practice, especially among informed and proactive patients. However, clinicians also identified key gaps, such as the absence of passive, avoidant, or digitally excluded patients. These findings reinforced the need to approach behavioral modelling as an iterative and interpretive process.

Clinicians described behavior not as a fixed trait but as a state that changes over time. Patients move across different archetypes depending on emotional resilience, treatment effects, disease progression, and external support. This dynamic nature requires systems that can adjust to change, not ones that rely on static classification. The eight patient archetypes identified in this chapter provide a practical framework for understanding this complexity. They also serve as design prompts for tailoring care strategies that match different motivations, needs, and communication styles.

The research questions guiding this chapter asked whether data-driven behavior models can align with clinical knowledge and how behavioral diversity can inform care design. The findings answered both.

Behavioral profiles can support clinicians, but only when grounded in practice. Behavioral diversity reveals clear paths for more personalized, adaptive systems. By identifying what patients value, how they make decisions, and when they change, designers can create tools that adjust care to meet patients where they are.

These insights move the discussion beyond behavior as data. They place behavior at the center of how care unfolds. Future systems must reflect this reality. They must observe change, support clinicians in understanding it, and help patients stay engaged through it. This chapter lays the foundation for such systems by offering a model of collaboration between technology and clinical insight one that treats behavior not as a category but as a signal to act.

Chapter 6
Systemic challenges of care pathways

This chapter shifts focus from disease-specific workflows to the broader forces that shape chronic care delivery. Drawing on insights from co-creation workshops with healthcare professionals and behavioral profiles generated from LLM analysis, we mapped how clinical, behavioral, and organizational elements influence the care journey. The goal of this chapter is to uncover why chronic care often fails to adapt to individual needs and where design interventions can have the most systemic impact. We tackle the following research questions in this chapter.

“What are the systemic influencers for a personalized care pathway?”

“How might we envision a future where outcomes become measurable and personalization becomes evidence based?”

The purpose of this chapter is threefold:

1. Reveal the current systemic state and show how clinical, behavioral, and organizational factors interact and sometimes clash.
2. Distil the principal system level problems that limit personalization and increase clinician workload.
3. Translate the problems into stakeholder needs and use them to frame a future state where outcome aligned feedback loops support adaptive care.

By the end of the chapter the reader will understand why many care pathways stall and what must change before design interventions can deliver truly personalized and feedback driven care.

6.1 Current State Influence map

An added benefit of the co-creation sessions was the opportunity to engage directly with healthcare professionals to better understand care pathways and their systemic challenges. Insights from these semi-structured interviews shaped the influence map (see fig 6.1), revealing the multifaceted dynamics of chronic care delivery.

At the center of this map lies the care pathway, influenced by both clinical factors such as disease progression and diagnostic complexity, and human elements, including communication styles, shared decision-making, and trust-building. While healthcare professionals aim to personalize care, they often work within the constraints of brief 15-minute consultations, relying heavily on clinical experience and instinct. This introduces a higher risk of misdiagnosis and limits the ability to systematically tailor care to each patient.

This current approach to personalization is time-consuming, clinician-centered, and varies significantly depending on the healthcare provider's experience. Despite the availability of clinical guidelines, the system lacks a standardized, evidence-based, and time-efficient method for personalization, one that also accounts for the patient's lived experience, preferences, and behavioral patterns.

The map also surfaces the behavioral diversity among chronic patients, captured through LLM-derived patient profiles validated by clinicians. These evolving behaviors significantly influence how patients engage with their care journey. Yet, clinical decisions continue to prioritize biomedical factors over human dimensions, such as patient education, trust, and quality of life.

Several systemic bottlenecks are apparent:

- Outcome measurement is inconsistent and lacks standardization.
- Goal-setting remains a challenge for many patients.
- Feedback loops that could support real-time care adaptation are underutilized.

Additionally, follow-ups are managed on a case-by-case basis, leading to increased time and burden on clinicians, higher costs for patients, and reduced system efficiency.

Overall, this influence map provides a structured yet human-centered lens on current care delivery. It highlights both the constraints and opportunities within chronic care, especially the potential of feedback-driven personalization and behavior-informed service design to improve adherence, reduce clinician burden, and ultimately enhance patient outcomes.

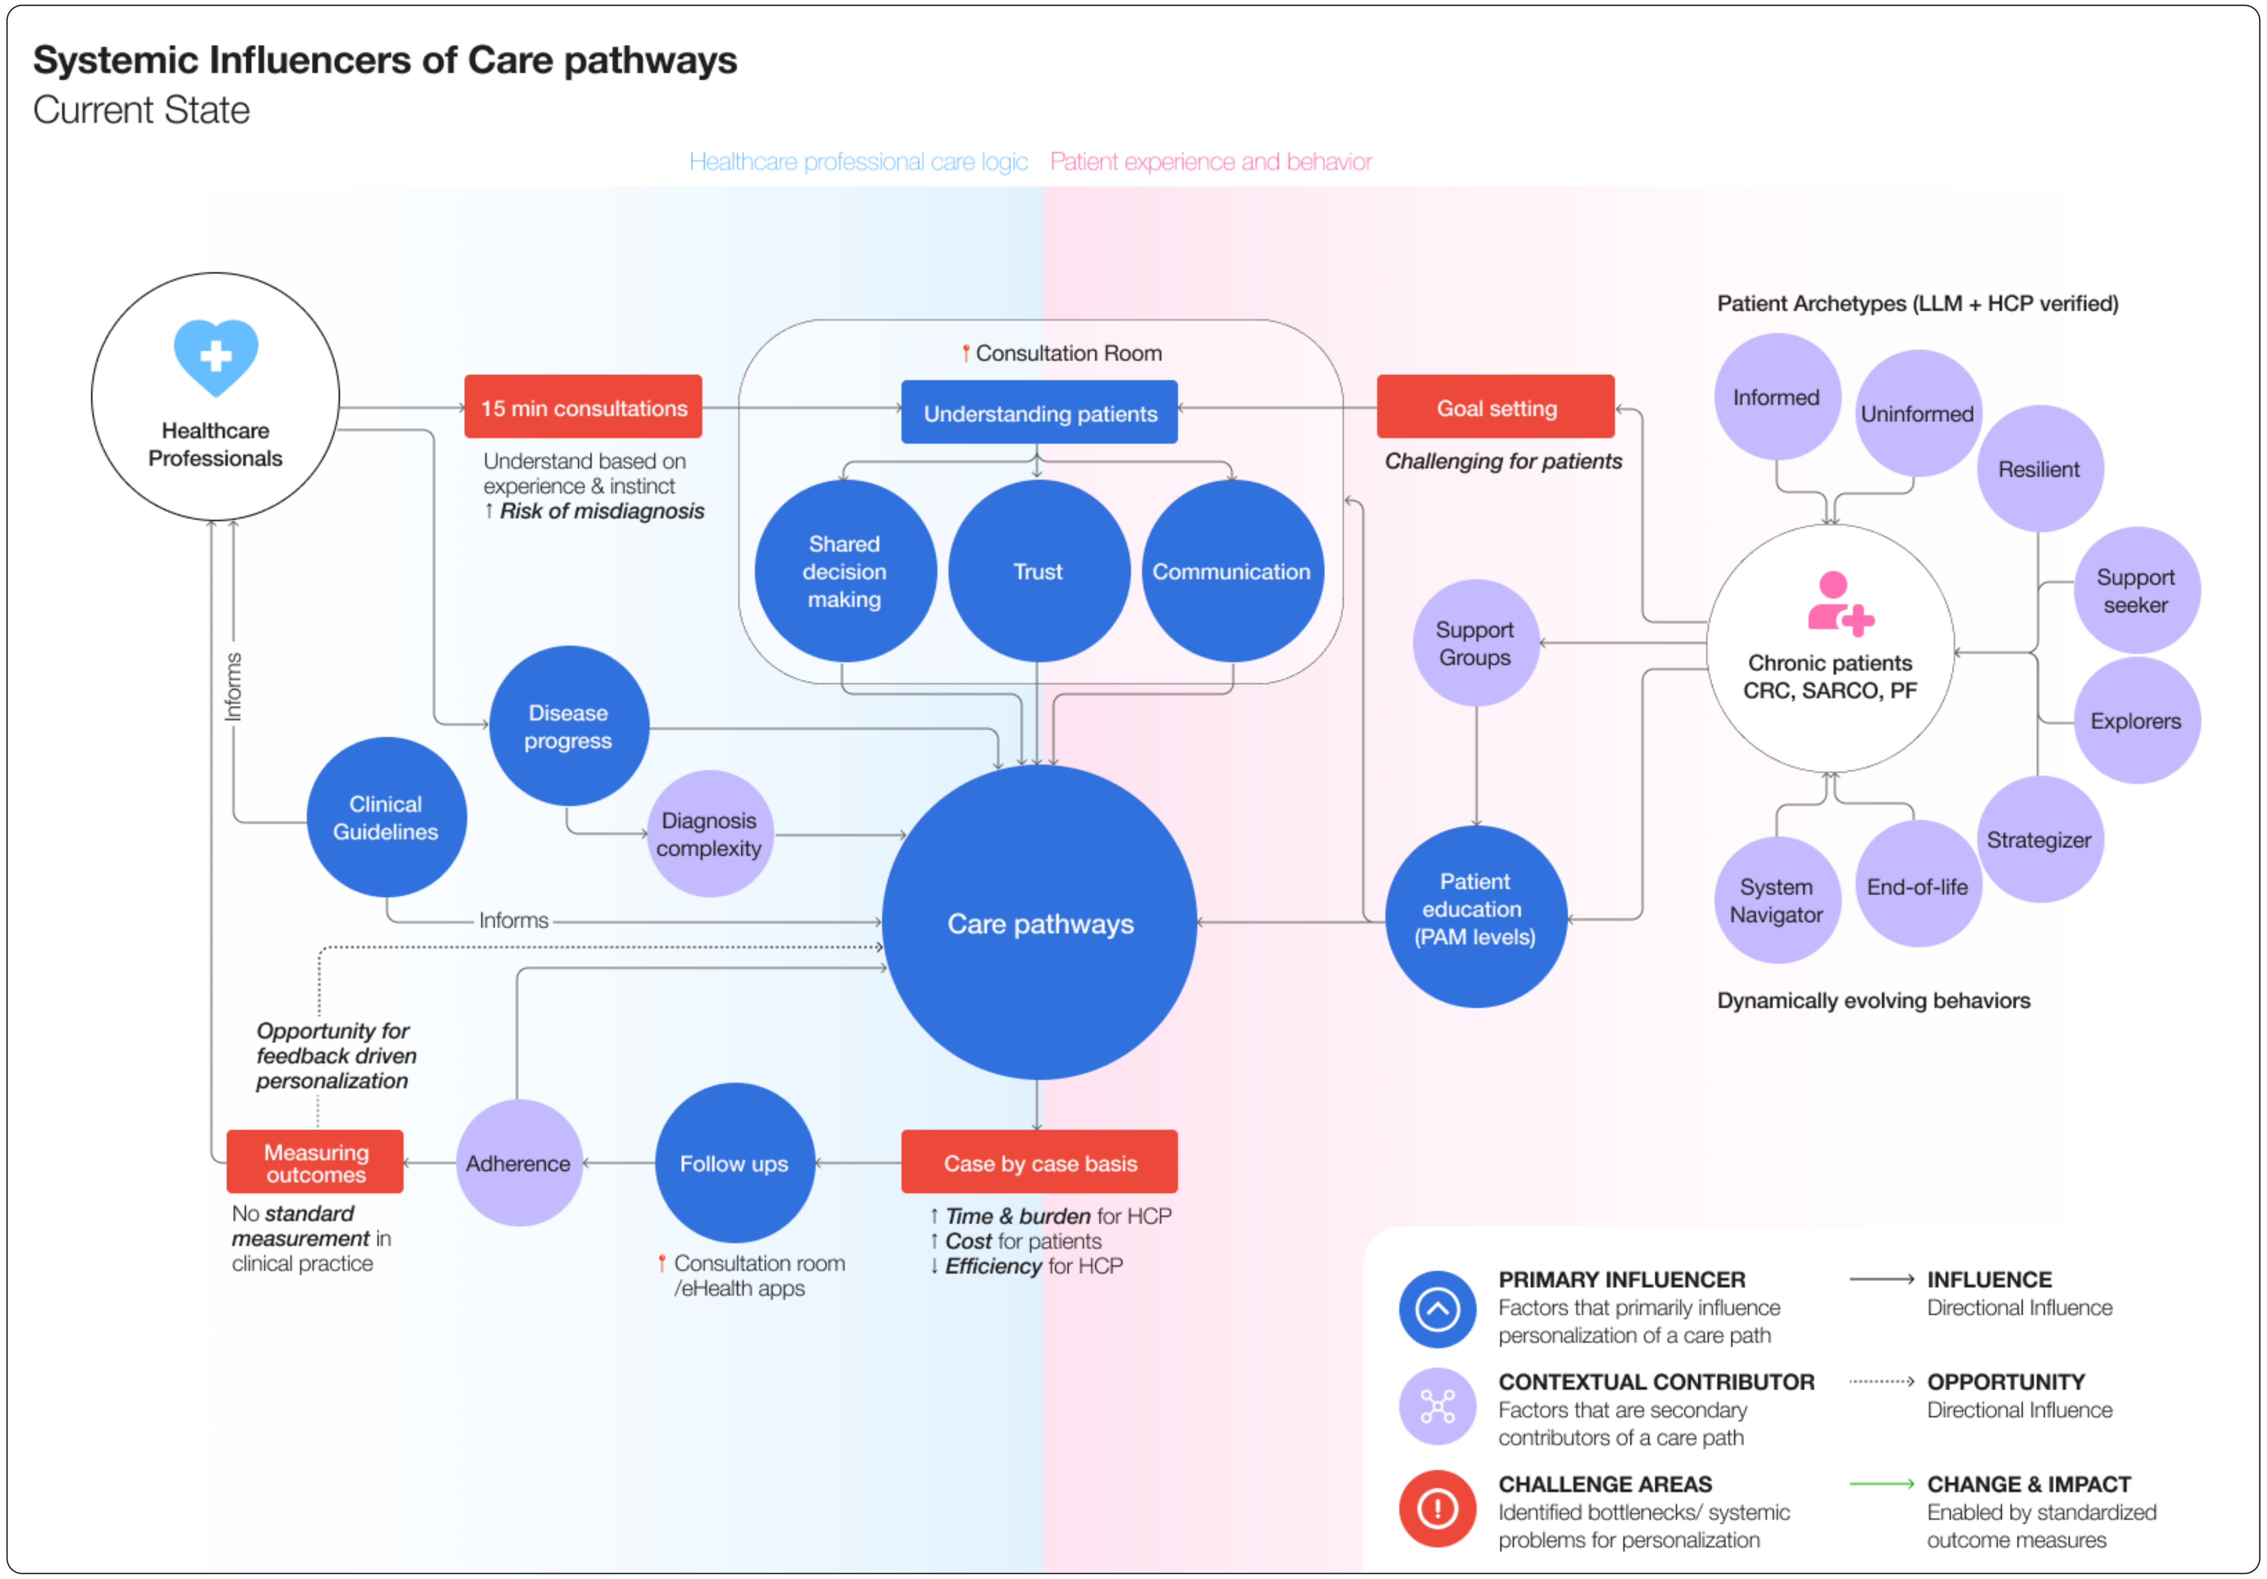

Fig 6.1 Current State - Systemic Influencers of Care pathways in chronic disease management

6.2 Envisioned future state

Building on the systemic challenges identified in the current care pathway, we recognized a pivotal opportunity that is, to shift from intuition-based personalization to feedback-driven, behavior-informed care. Informed by healthcare professionals' lived realities and constraints, our investigation revealed one transformative enabler: the integration of standardized outcome metrics into daily clinical workflows.

The resulting "future state" influence map (see fig 6.2) illustrates a reimagined chronic-care system that is anchored in data, adaptable to patient behavior, and designed to reduce systemic inefficiencies.

Moving from experience to evidence

In this future system, the care pathway remains central but is now dynamically informed by real-time outcome data. This enables a continuous evidence loop, where insights from measured outcomes feed directly into disease tracking and reduce diagnostic complexity. Uncertainty is gradually replaced with clarity, reducing clinician dependence on experience alone.
Within the consultation room, goal setting evolves from an isolated task to a feedback-driven process, guided by individual outcome trends. This allows healthcare professionals to anchor even short 15-minute consultations in data, improving confidence in decision-making. For patients, this results in clearer expectations, more achievable goals, and a stronger sense of involvement in their care.

Making personalization scalable

LLM-derived behavioral profiles still serve as the backbone for personalization, but two major gaps are now addressed:

1. Adherence becomes visible and trackable through standardized outcome metrics.
2. Follow-ups transition from ad hoc to benchmark-basis, ensuring consistency and reducing the burden on both clinicians and patients.

This change is reflected in the green arrows on the map, signaling key improvements such as:

- Decreased diagnostic complexity
- Measurable and adaptive adherence support
- Reduction in clinician workload
- A shift toward personalization cues (e.g., trust-building, communication style) that match patient literacy levels and behavioral profiles.

Previously identified red bottlenecks such as outcome measurement and goal setting are now transformed into green solution points. What was once a missing component becomes the core engine of continuous adaptation. The redesigned process empowers both clinician and patient to co-create care plans grounded in real data and personalized behavioral understanding.

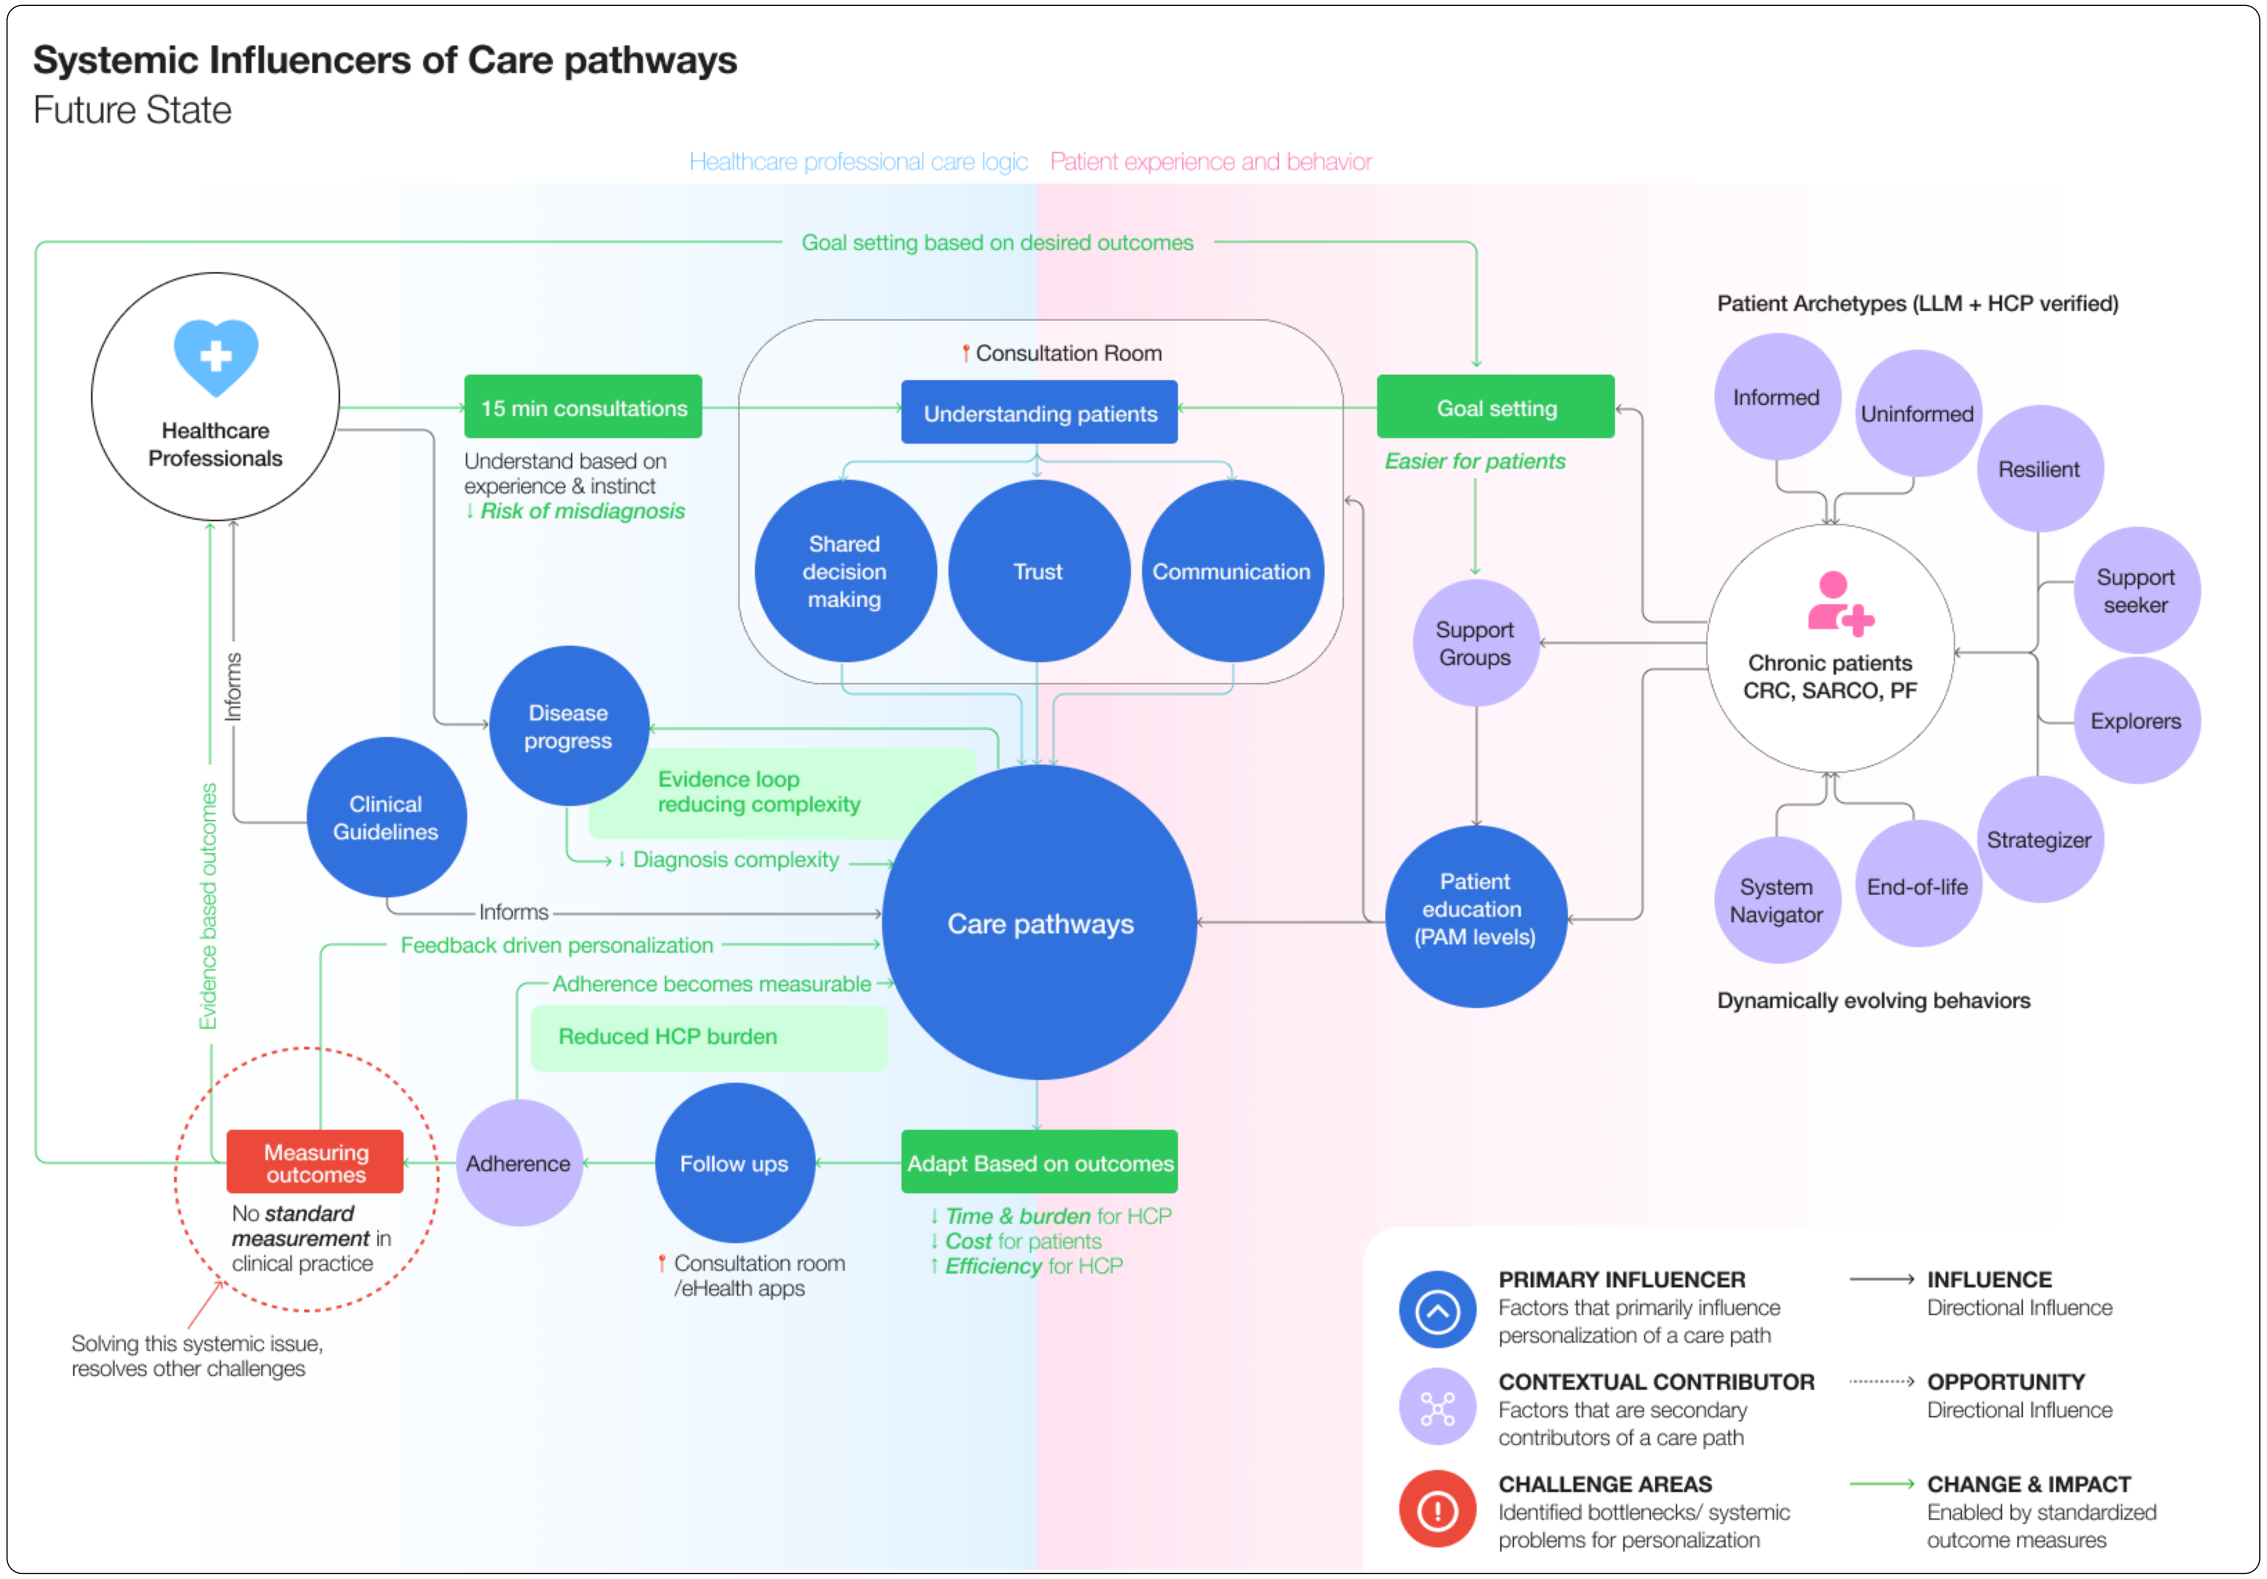

Fig 6.2 Future state - Systemic Influencers of Care pathways in chronic disease management

6.3 Vision 2040

The next fifteen years offer a pivotal opportunity to address the growing complexity of chronic care. During this period, we can move towards a system that is adaptive, evidence based and centered around both clinical outcomes and patient experiences. By 2040, we envision a model of chronic care in which each care pathway evolves continuously, guided by real time feedback, insights from patient behavior, and standardized outcome metrics that inform decisions for both clinicians and patients.

The selection of 2040 as the target year is intentional and supported by multiple converging trends.

1. The global burden of chronic disease is mounting rapidly. The World Health Organization's latest fact sheet (updated December 2024) shows that noncommunicable diseases (NCDs) were responsible for at least 43 million deaths in 2021, accounting for approximately 75 percent of all non-pandemic deaths worldwide (World Health Organization: WHO, 2024b). Building on these trends, the Global Burden of Disease 2021 forecasting analysis released by the Institute for Health Metrics and Evaluation in The Lancet (May 2024) projects that by 2040, NCDs will account for well over three quarters of all global deaths, approaching four in every five households (Vollset et al., 2024). Taken together, these updated estimates highlight the urgency of reengineering health systems toward long term prevention and management of chronic conditions.

2. The technological foundation for delivering adaptive care is becoming more robust. The adoption of digital health infrastructure including electronic health records, artificial intelligence assisted decision support, and patient engagement platforms is advancing rapidly. According to the Organization for Economic Co-operation and Development, digital health strategies across member countries, including founding members such as the Netherlands, are expected to reach broad maturity within the next ten to fifteen years (Health at a Glance 2023).

3. By 2040, a new generation of digitally fluent patients and clinicians will be active within the healthcare system. These users will be more comfortable with digital tools, behavior tracking, and data driven personalization, enabling deeper integration of human factors into clinical routines (Javaid et al., 2024; Bucher, 2023). This transformation is driven by three major shifts in how chronic care is delivered and experienced.

This transformation is driven by three major shifts in how chronic care is delivered and experienced.

1. **From reactive to adaptive care**
   Current care models often respond to clinical events after they occur. In contrast, adaptive care enables proactive intervention. This is made possible by embedding real time outcome tracking such as symptom monitoring and behavioral feedback, into everyday clinical practice. Early success with remote monitoring technologies has already shown reductions in emergency visits and improved treatment adjustments. Standardized outcome measurement allows care pathways to become responsive and precise.

2. **From intuition to evidence informed decision making**
   As seen in the earlier chapters, under time pressure, many clinicians rely heavily on instinct and experience. While this is often necessary, it introduces variation and increases the risk of misdiagnosis. By 2040, we anticipate the widespread use of structured patient data, behavior-based profiles, and decision support tools. These will allow clinicians to base their recommendations on both longitudinal evidence and real time metrics. Even short consultations can become more effective when guided by standardized frameworks like the Patient Activation Measure and PROMIS tools (Patient-Reported Outcomes Measurement Information System | NIH Common Fund, n.d.).

3. **From clinician directed to co-created care**
   At present, care personalization is shaped largely by the clinician's perspective. However, behavioral profiling from LLM analysis in the previous chapters, reveals that patients engage with care in highly diverse ways. Some seek information actively, others rely on support and reassurance. By 2040, systems will be capable of tailoring communication and decision-making approaches based on patient readiness, comprehension, and emotional needs. This will enable more meaningful participation and improve trust, adherence, and health outcomes (SolvEdge, n.d.).

   Together, these shifts will reframe care pathways as intelligent and evolving systems. They will learn from each individual and adapt over time, leading to faster diagnoses, better treatment alignment, and measurable improvements in both patient quality of life and overall system efficiency.

By correlating the multiple converging trends, and envisioning a future where human metrics would guide care decisions, I came up with a vision for personalizing chronic care by the year 2040.

> By 2040, care pathways will use human metrics like emotion, motivation, and behavior to guide decisions. Care will adapt in real time, shaped by patient behavior and built together with clinicians for better and more personal outcomes.
>
> *- Vision 2040*

This vision creates the foundation for a structured transformation plan for personalized chronic care. The strategic roadmap outlined in Section 6.4 details how this future can be achieved, starting from targeted pilots and advancing toward integrated system wide implementation.

6.4 Strategic Roadmap Turning Vision 2040 Into Action

In the previous section, we set out an ambition for chronic-care pathways that learn continuously from clinical outcomes and human factors such as motivation and trust. This section converts that ambition into a time-bounded sequence of steps in the form of a strategic roadmap.The roadmap spans fifteen years and is organized into three horizons (see fig 6.3).

The first horizon (2025–2027) corresponds to a phase of focused piloting within a single academic medical center, where early validation and feedback are crucial. This short time frame supports rapid iteration while minimizing complexity and risk.

The second horizon (2027–2034) aligns with national and international digital health infrastructure timelines, including the anticipated maturation of interoperable EHRs and wearable health tech integration across the Netherlands (Health at a Glance, 2023). This seven-year period allows sufficient time for scale-up, multi-site deployment, and integration of wearable and behavioral data into routine workflows.

The third horizon (2035–2040) captures the long-term goal of embedding behavioral AI, human metrics, and adaptive decision-making into national policy and practice. It aligns with broader European policy developments, including full enforcement of the EU AI Act and the European Health Data Space, expected to standardize governance and auditability by the late 2030s.

Each horizon is further discussed in detail, capturing several dimensions starting with strategic goals, emerging trends, technology evolution, behavioral touchpoints, business models, and ethical safeguards. Together, they outline a clear and adaptable path for transforming chronic care from its current reactive form into a proactive, evidence-informed, and co-created system by 2040.

6.4.1. Horizon 1

Build and Prove (2025 to 2027)

Goals

We aim to launch a pilot at Erasmus MC using a basic patient app and clinician-facing web plugin. Around 300 patients (100 patients each from Colorectal cancer, sarcoidosis and pulmonary fibrosis) will test the platform. This limited scope allows for controlled testing and quick iterations.

Key trends

Clinicians are increasingly aware of behavioral differences among patients and want tools to personalize care. Digital health experiments are growing, particularly those that focus on emotional and behavioral data. Regulatory bodies, including the European Commission, allow decision-support software to be fast-tracked if it does not automate clinical decisions and ensures human oversight (Medical Device Coordination Group, 2021).

Value created

Patients are guided to notice and journal stress, mood, and confidence, which reduces emotional overload and sharpens self-management. In this way, patients are better equipped to handle stress and emotional burden. At the same time, clinicians gain structured insights into behavior, which were previously missing from standard visits.

Outcomes and KPIs

1. At least sixty percent of users engage with the app on five or more days each week. Meta-analyses of digital monitoring show that adherence above sixty percent is strongly associated with clinical benefit (Po et al., 2024).

2. Mean Patient Activation Measure increases by one point. Even small increases in PAM are linked to better health outcomes and fewer hospitalizations (Abdelraheem et al., 2024).

3. We aim for 2x more improvement in clinician-perceived usefulness of behavioral data ensuring that insights are clinically meaningful and used in care decisions.

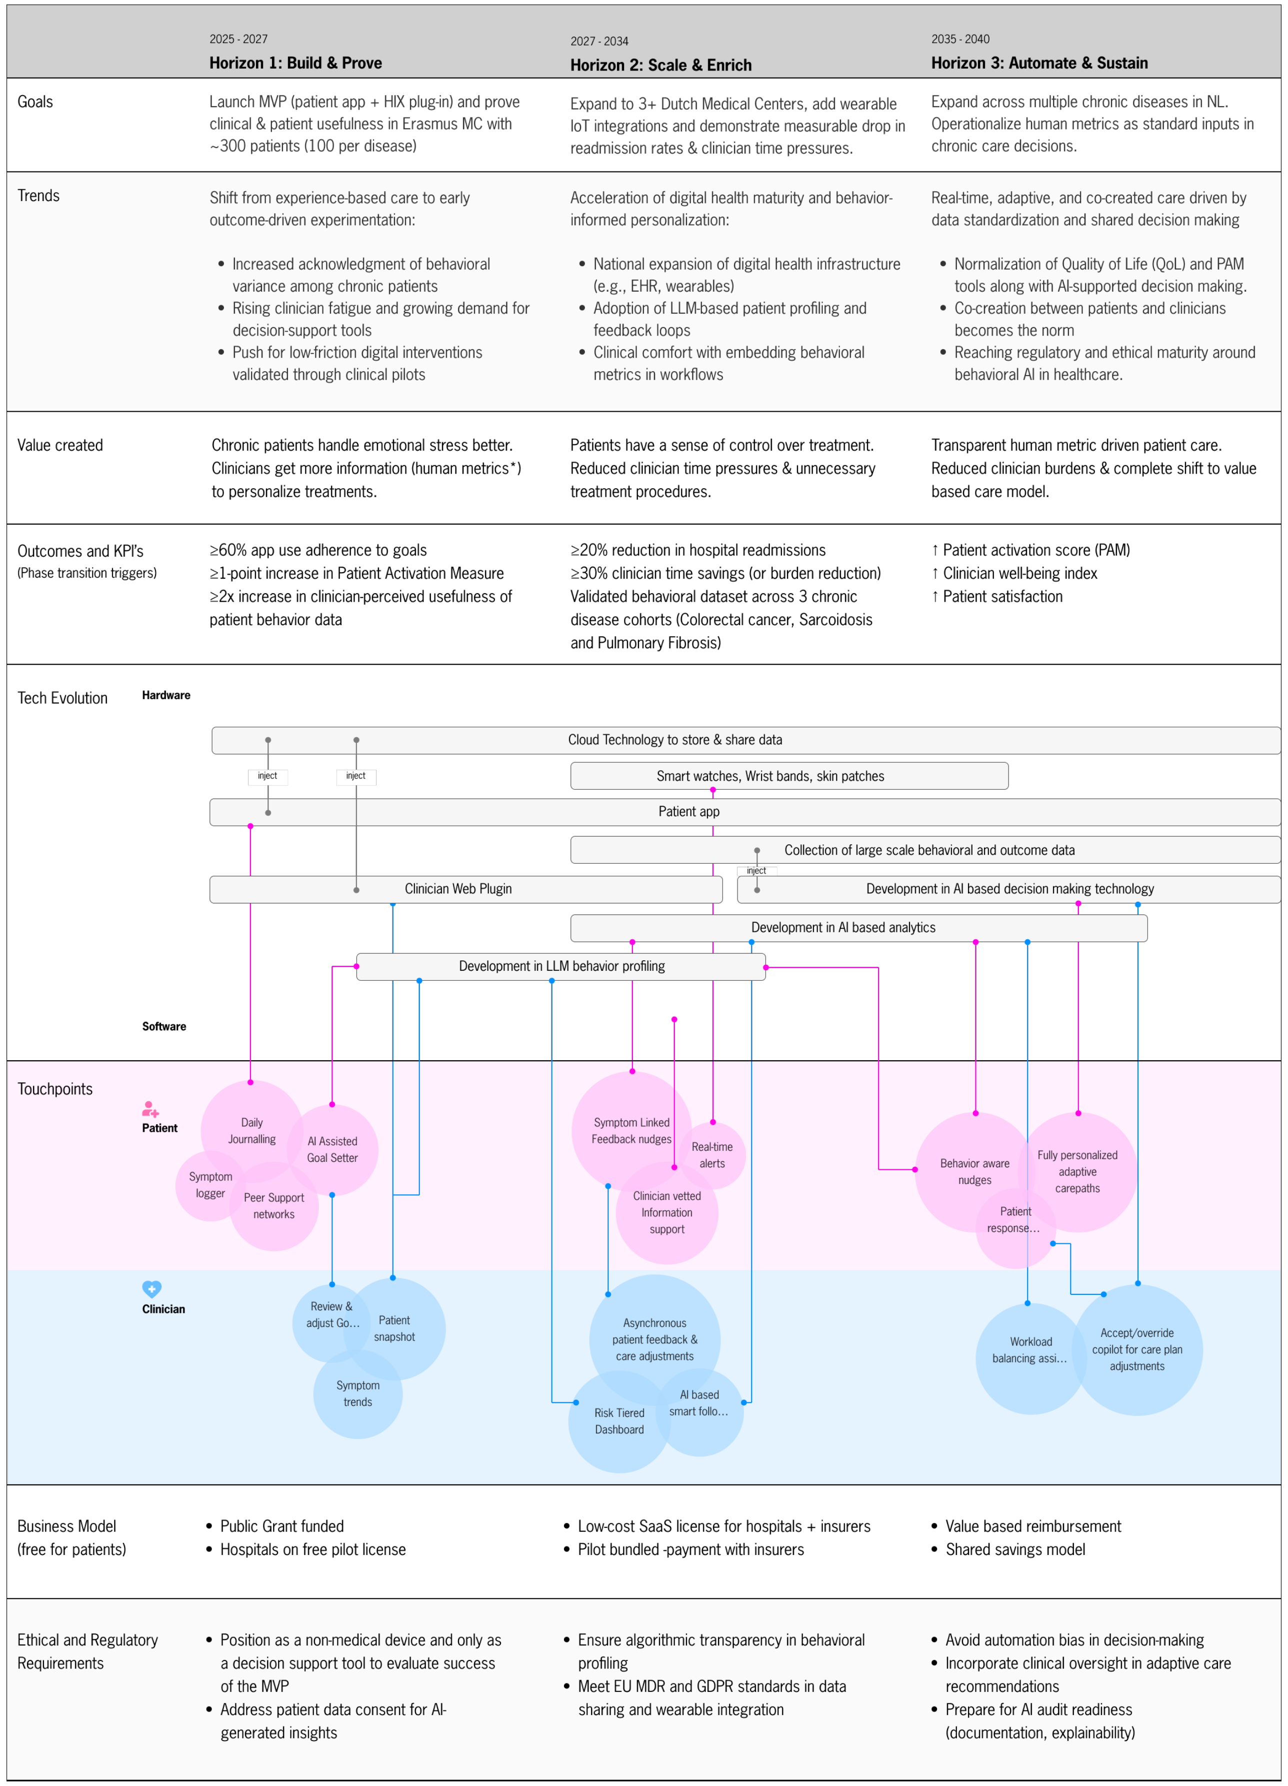

	2025 - 2027 **Horizon 1: Build & Prove**	2027 - 2034 **Horizon 2: Scale & Enrich**	2035 - 2040 **Horizon 3: Automate & Sustain**
Goals	Launch MVP (patient app + HIX plug-in) and prove clinical & patient usefulness in Erasmus MC with ~300 patients (100 per disease)	Expand to 3+ Dutch Medical Centers, add wearable IoT integrations and demonstrate measurable drop in readmission rates & clinician time pressures.	Expand across multiple chronic diseases in NL. Operationalize human metrics as standard inputs in chronic care decisions.
Trends	Shift from experience-based care to early outcome-driven experimentation: • Increased acknowledgment of behavioral variance among chronic patients • Rising clinician fatigue and growing demand for decision-support tools • Push for low-friction digital interventions validated through clinical pilots	Acceleration of digital health maturity and behavior-informed personalization: • National expansion of digital health infrastructure (e.g., EHR, wearables) • Adoption of LLM-based patient profiling and feedback loops • Clinical comfort with embedding behavioral metrics in workflows	Real-time, adaptive, and co-created care driven by data standardization and shared decision making • Normalization of Quality of Life (QoL) and PAM tools along with AI-supported decision making. • Co-creation between patients and clinicians becomes the norm • Reaching regulatory and ethical maturity around behavioral AI in healthcare.
Value created	Chronic patients handle emotional stress better. Clinicians get more information (human metrics*) to personalize treatments.	Patients have a sense of control over treatment. Reduced clinician time pressures & unnecessary treatment procedures.	Transparent human metric driven patient care. Reduced clinician burdens & complete shift to value based care model.
Outcomes and KPI's (Phase transition triggers)	≥60% app use adherence to goals ≥1-point increase in Patient Activation Measure ≥2x increase in clinician-perceived usefulness of patient behavior data	≥20% reduction in hospital readmissions ≥30% clinician time savings (or burden reduction) Validated behavioral dataset across 3 chronic disease cohorts (Colorectal cancer, Sarcoidosis and Pulmonary Fibrosis)	↑ Patient activation score (PAM) ↑ Clinician well-being index ↑ Patient satisfaction
Business Model (free for patients)	• Public Grant funded • Hospitals on free pilot license	• Low-cost SaaS license for hospitals + insurers • Pilot bundled -payment with insurers	• Value based reimbursement • Shared savings model
Ethical and Regulatory Requirements	• Position as a non-medical device and only as a decision support tool to evaluate success of the MVP • Address patient data consent for AI-generated insights	• Ensure algorithmic transparency in behavioral profiling • Meet EU MDR and GDPR standards in data sharing and wearable integration	• Avoid automation bias in decision-making • Incorporate clinical oversight in adaptive care recommendations • Prepare for AI audit readiness (documentation, explainability)

*Human metrics include factors such as emotional health including mood, stress, motivation, confidence and social well-being factors

Fig 6.3 Reimagining Care Pathways
A behavioral and clinical Integration roadmap

Key Stakeholders
Patients
Clinicians
Hospitals
Insurers

Technology evolution
Only smartphone-based logging and cloud storage are used. Wearables are excluded at this stage to minimize complexity and privacy risks.

Touchpoints
Patients contribute through daily journals, receive AI-assisted goal-setter prompts, and can join peer-support networks. Clinicians are provided with a dashboard summarizing symptom trends and receive patient progress snapshots before consultations.

Business model and governance
A government funded grant can be used for initiating and testing the pilot. Software is positioned as a non-medical decision-support tool to simplify regulatory clearance.

6.4.2. Horizon 2
Scale and Enrich (2027 to 2034)

Goals
After meeting the initial KPIs, the adaptive-care platform will scale to three Dutch academic medical centers. During this phase, wearable devices such as smartwatches, wristbands, and skin patches will be integrated, enabling continuous collection of physiological and activity data. This data will flow directly into electronic health records and be combined with hospital operations metrics to evaluate reductions in hospital readmissions and clinician workload across varied patient groups.

Emerging trends
National health systems are investing in interoperable digital infrastructure (Health at a Glance 2023). In this time period, Large-language-model pipelines will mature, enabling richer patient profiling and automated feedback summaries that cut documentation time. A 2024 Lancet Digital Health evaluation demonstrated that LLM-assisted triage notes significantly reduced clerical workload without harming safety (S. Chen et al., 2024). These factors will contribute to the rise in clinician interest and behavioral feedback loops.

Value created
In this horizon, we can predict that patients gain a clearer sense of control over treatment because their own wearable data would validate goal progress in near real time. Clinicians benefit from reduced burden through smarter triaging and fewer routine follow-ups.

Outcomes and KPI's
With the integration of behavior-based feedback loops, we anticipate hospital readmissions to fall by at least twenty percent. With the introduction of clinician facing dashboards, we would see a reduction in routine consultation time by directing clinician focus towards shared decision making that meets patient goals. Another important result from the previous horizon carried forward to horizon 2 is the validation of behavioral datasets across the three chronic conditions (colorectal cancer, sarcoidosis, pulmonary fibrosis). This would further enable development of accurate predictive models that can further be scaled to other conditions.

Technology evolution
Wearable data is merged with behavioral and clinical records. Data pipelines standardize ingestion of step counts, heart-rate variability, and sleep metrics for integration with outcome dashboards.
Development begins on explainable-AI analytics to meet the upcoming EU Artificial Intelligence Act requirements for transparency in high-risk medical systems (Artificial Intelligence in Healthcare, 2025).

Touchpoints
Patients receive symptom-linked behavior nudges and real-time alerts about their condition and changes in treatment plans. Clinicians move towards asynchronous review using a risk-tiered dashboard and AI-based smart tools.

Business model
Hospitals license the platform through a low-cost software-as-a-service contract, while insurers test bundled-payment arrangements that share savings generated by prevented admissions, an approach already explored in Dutch bundled-care pilots (Bour et al., 2023; Karimi et al., 2021).

Ethical and regulatory safeguards
Algorithmic transparency reports are shared with both clinicians and patients. All systems must meet GDPR and data exchanges must conform to the upcoming European Health Data Space principles of citizen control and secondary-use governance (European Health Data Space Regulation (EHDS), 2025).

6.4.3. Horizon 3
Automate and Sustain (2035 to 2040)

Goals
The goal for Horizon 3 is for the adaptive-care platform to reach nationwide coverage across multiple chronic-disease domains. Human metrics such as confidence levels and emotion trends shall become standard inputs to decision-support algorithms, fully realising the Vision 2040 statement.

Key trends
With successful integration of behavior data into clinical workflows, we can anticipate a future with real-time, adaptive and co-created care driven by data standardization and shared decision making. Quality of Life metrics and PAM scores would turn into mainstream clinical vital signs giving teams a common language for symptoms, function, and readiness to self-manage.

With advanced access to behavior data from previous horizons, Large-language-model copilots would have the ability to generate personalized options and highlight trade-offs in plain language. Professional guidance would insist on explainable outputs that clinicians and patients can verify before acting, reducing mistrust and alert fatigue.

Value created
Care pathways adjust automatically, offering behavior-aware nudges and personalized care packages. Clinician burnout lessens as cognitive load shifts to verifiable AI. Patients perceive a transparent partnership grounded in shared data.

Outcomes and KPIs
1. Sustained improvement in patient-activation scores confirms that co-created decision making continues to motivate self-management.
2. Clinician-wellbeing index scores rise, tracking reduced administrative burden.
3. Patient-satisfaction ratings climb, confirming that adaptive care meets experiential expectations.

These KPIs correspond to the quadruple-aim framework, which emphasizes patient experience, population health, cost containment, and provider wellbeing.

Technology evolution
AI generates personalized care plans and updates them continuously. Bias detection and explainability are built in, meeting audit requirements under the finalized EU AI Act (Regulation - EU - 2024/1689 - EN - EUR-LEX, n.d.).

Business and ethical requirements
Shared-savings agreements between hospitals and insurers will become mainstream because data now quantify reductions in low-value procedures and unnecessary diagnostics. A national body or an organization would then be required to certify the AI modules and enforce policies that preserve human authority over final decisions.

To summarize, this roadmap provides a clear and practical path from pilot to nationwide transformation. Each horizon adds scale, sophistication, and safety while ensuring patients and clinicians remain at the center. KPIs are selected not just to show progress, but to confirm value across system layers from individual outcomes to economic sustainability. By 2040, chronic care can evolve into an adaptive, human-centered system that learns from every patient it serves.

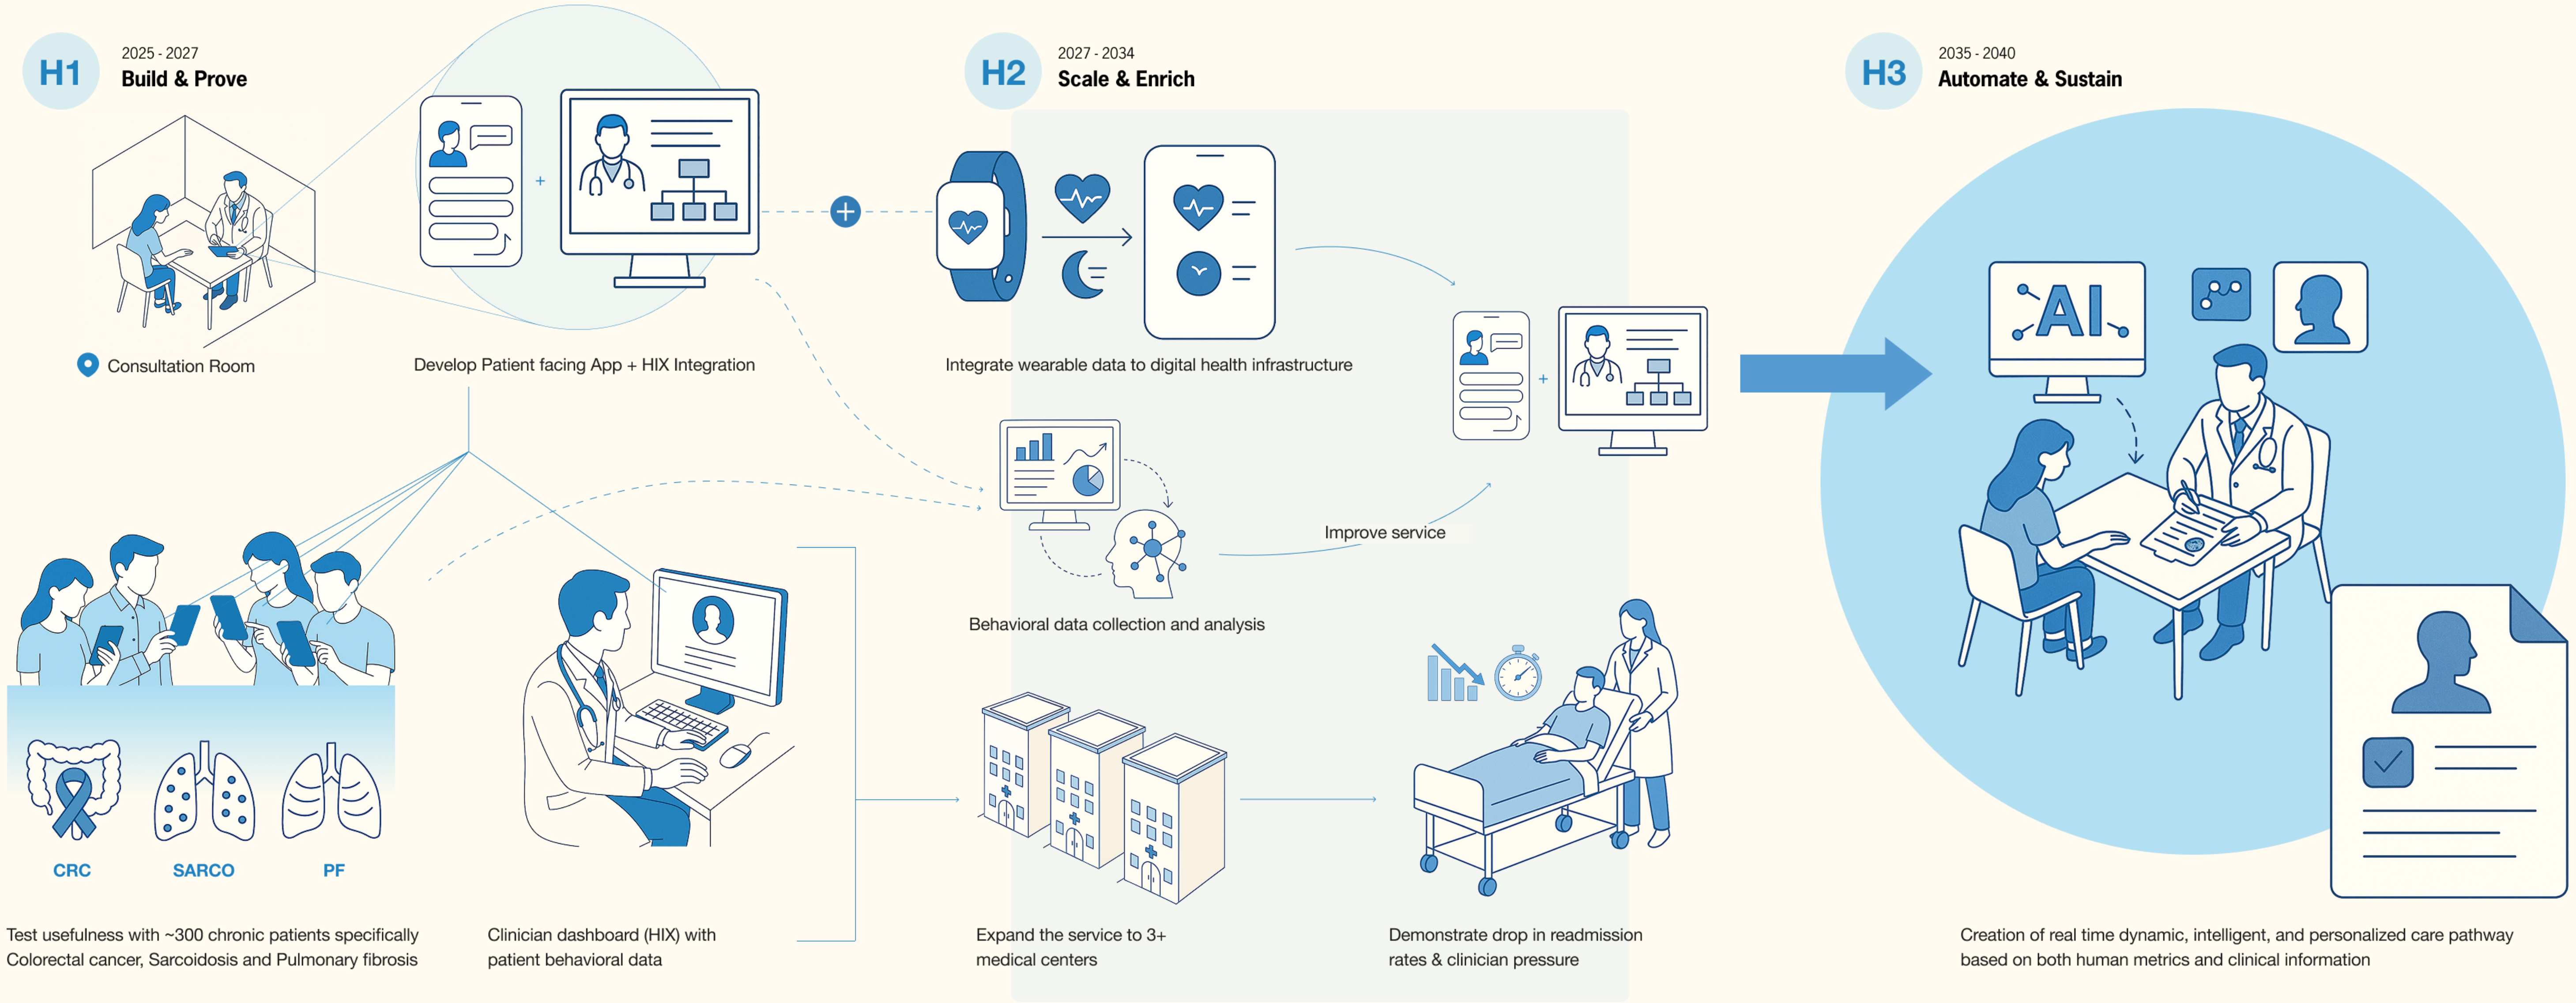

Fig 6.4 Visual journey of implementation from H1- H3

6.5 Reflections

Taking a step back and engaging with both the systemic realities and future possibilities of chronic care delivery revealed how much of healthcare is shaped by invisible constraints. Time pressure, fragmented feedback, and behavioral signals often go unnoticed. While these barriers are well known to clinicians, mapping them out as interlocking systems made their weight and complexity far more tangible.

One of the most personally challenging yet rewarding parts of this process was learning how to build a strategic roadmap. Road mapping was entirely new to me. At first, it felt abstract, balancing vision, feasibility, and timing across multiple layers of a healthcare system was not something I had done before.

But as I revisited the previous chapters, grounding myself in the challenges faced by clinicians, I was able to connect the dots between human metrics (communication, trust, behavior etc.) and system gaps. This helped me understand how transformation unfolds in real healthcare settings and that care is not just about decisions, but about timing, trust, and presence.

Ultimately, this chapter taught me that transforming chronic care is not just a matter of adding new technologies. It requires rethinking the rhythm and feedback of care itself. The strategic roadmap became a tool for seeing clearly how systems evolve, where trust is built, and why personalization must be woven into care delivery.

6.6 Conclusions

This chapter traced the complex, often fragmented landscape of chronic care and identified a path toward transformation grounded in both clinical realism and behavioral intelligence. The influence maps revealed how personalization efforts are often limited not by lack of intent but by the absence of measurable feedback, behavioral clarity, and scalable tools.

We were able to define a clear and compelling vision for 2040 - A care system that adapts in real time, guided by patient-reported outcomes and co-created decisions. The strategic roadmap translated that vision into actionable steps, balancing feasibility and ambition. Each horizon was designed to build clinical trust and behavioral relevance along with technical capability.

Together, these insights reinforce a critical shift in thinking, away from static care models and towards living systems of care that learn, respond, and grow with each patient. They also highlight the role of design as a strategic bridge between system constraints and human needs.

These foundations now guide the next chapter. In the next chapter, we move from planning to building a design concept by translating the strategic intent of roadmap into real-world touchpoints that support adaptive, behavior-aware chronic care.

Chapter 7
Moving towards a solution

In the previous chapter, we explored the systemic barriers that limit personalization in chronic care and proposed a roadmap toward a more adaptive and feedback-driven future. We identified clinical, behavioral, and organizational factors that influence current care delivery and framed a vision for 2040 where real-time human metrics support decisions and care pathways evolve continuously. However, to move beyond strategic intent, we needed to explore how this transformation might begin to take shape in the lived experience of patients.

This chapter shifts from vision to execution. It focuses on the conceptual development of patient touchpoints that can be used for capturing behavioral data as outlined in the strategic roadmap. These design explorations help uncover what kinds of interactions, data flows, and emotional cues are needed to support self-management in a personalized care system. While these features have not been tested with users, they serve as critical building blocks for imagining how future tools might sense, respond to, and adapt around patient behavior. We investigate the following research questions:

“How can emotionally grounded patient pain points be translated into actionable digital features that support self-management?”

“In what ways can behavioral data be meaningfully captured through everyday patient touchpoints without adding emotional burden?”

“How might personalized digital flows improve patient engagement and support shared decision making in care journeys?”

“What role can clinician-facing dashboards play in transforming emotional and behavioral insights into actionable care adjustments?”

I walk through the full design process, beginning with insights from large language model analysis and co-creation sessions that revealed patient’s pain points. These insights informed a round of early ideation that shaped the broader solution space. I then describe how we defined a focused feature set, created user flows, and developed the interaction logic. The chapter ends with personal reflections on moments of creative blocks during the design process and how I managed to move forward.

7.1 Revisiting user pain points from LLM analysis

For patients living with chronic illnesses like colorectal cancer, sarcoidosis, or pulmonary fibrosis, the journey through care brings more than just physical symptoms. Emotional stress sits at the center of their experience. It shapes how they interpret symptoms, respond to treatment, and navigate their daily lives.

This emotional stress builds over time. Patients face repeated hospital visits, unpredictable test results, and constant shifts in treatment plans. These disruptions create a sense of fear and uncertainty. Many patients describe living in a state of tension, waiting for symptoms to return or treatment to change unexpectedly. This pressure often triggers anxiety, particularly around medical updates and appointments.

Side effects such as fatigue, discomfort, and hair loss create daily challenges. These symptoms reduce energy, alter body image, and affect confidence. Patients struggle to manage these changes while trying to maintain a sense of normalcy. Some hesitate to talk about their struggles, fearing judgment or misunderstanding. As a result, they begin to withdraw from support systems and isolate themselves.

The emotional burden grows heavier when patients try to understand their medical situation. Many feel overwhelmed by the volume of information they receive. They try to make sense of complex medical advice while also dealing with misinformation and conflicting messages online. They find it hard to trust sources, and even harder to make confident treatment decisions.

Patients often carry the weight of decision-making alone. They must explore treatment options, assess risks, and evaluate alternatives without clear guidance. This responsibility leads to decision fatigue. Many worry about choosing the wrong path, which causes hesitation and delays in care.

Lifestyle changes add another layer of difficulty. Patients adapt to new diets, altered work routines, and shifting family dynamics. These changes affect identity and autonomy. As they try to balance medical needs with personal responsibilities, many feel emotionally drained and unsupported.
Each of these pain points connects back to emotional stress. It sits at the root of the experience and influences every choice, thought, and reaction. When emotional stress rises, patients lose clarity. They struggle to engage with care plans, withdraw from communication, and lose trust in the process.
This pain point map (fig 7.1) shows how deeply emotional stress links to physical symptoms, decision making, and information overload. It reinforces the need to design care tools that address these emotional experiences directly.

Fig 7.1. Interlinked patient pain points identified from LLM analysis and clinician inputs.

7.2 Design activities

To begin the design activities, I grounded the process in patient pain points identified through LLM analysis and clinician input. I used a mix of reflective, creative, and structured design methods to translate these insights into actionable ideas.

7.2.1 Framing the right questions in the form of How Might We..?

I started by shaping the problem space through a set of focused How Might We (HMW) questions. These questions helped transform pain points into design directions:

- How might we reduce the emotional burden on the patient?
- How might we help patients navigate treatment options better?
- How might we help patients ease into their new lifestyle?
- How might we capture enough data for personalizing care?
- How might we deliver behavioral snapshots to clinicians?

Each question opened possibilities without narrowing down too quickly. They allowed me to keep user needs in focus while exploring varied design solutions.

7.2.2 Two-Step Ideation process Human First, AI-Enriched

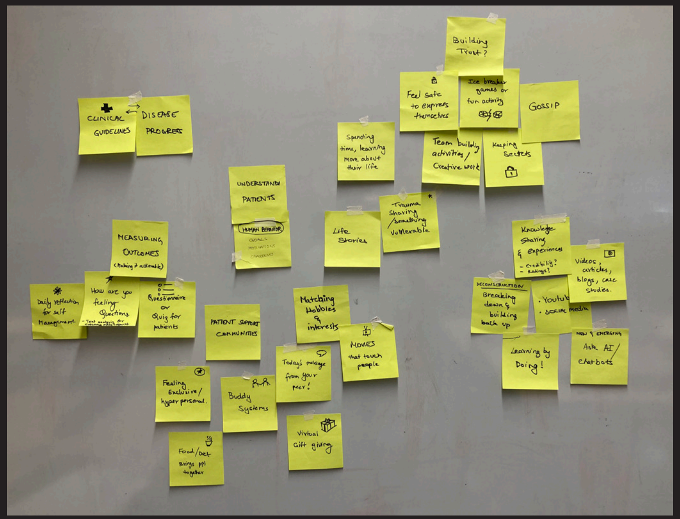
Fig 7.2. Brainstorming idea directions

I began ideating by writing down my ideas on post its, focusing on quantity. After this, I used ChatGPT to expand my ideas. I asked for examples, prompted alternatives, and explored new interaction models. I prompted the problem space I was trying to address and provided the ideas I came up with. I then prompted to expand my ideas focusing once again on quantity. This enriched the creative process. For instance, the idea of building trust evolved into small, playful elements like creative icebreakers or messages from a virtual care guide. While I generated the direction through human empathy, AI helped widen the design space and bring depth to early concepts.

7.2.3. Creating an Interaction Vision

To truly relate to the way patients feel and design successful interventions, I developed a guiding metaphor “***Lost in Space***”. A chronic diagnosis often leaves patients feeling unanchored, as if they have entered an unfamiliar world. They struggle to understand what is happening, whom to trust, and how to regain control.


Fig 7.3. Visualizing the metaphor of “Lost in space”

In this metaphor, the digital experience that I intend to design would act as navigation support. Healthcare professionals are positioned as ground control, offering timely, reassuring guidance. Patients receive cues, nudges, and checkpoints to help them move forward with confidence. This vision shaped the overall tone of the concept. It guided interaction principles such as clarity, warmth, and progressive onboarding. It also informed key ideas and behavior-linked prompts that respond to the patient’s emotional state.

7.2.4. Why a digital app experience?

I chose to develop a digital application because it fits the daily routines of most patients. People already rely on mobile devices to manage appointments, record symptoms, and seek information. A digital format allows continuous, flexible interaction and supports asynchronous communication with clinicians.

More importantly, a digital product creates the opportunity to collect behavioral data over time. This information can help track mood, motivation, and readiness, which are central to building personalized care. The digital layer enables both patient-facing tools for reflection and Clinician-facing dashboards, making it the ideal format for capturing and presenting information.

This design phase turned high-level insights into concept-ready features. Through targeted ideation, metaphor-driven framing, and focused reasoning, I built a foundation that connects patient needs with interaction possibilities. In the next section, I define a set of feature cards. Each card will describe the user value, desired behaviors, user actions, supporting data points, and the mechanisms that make them work.

7.3 Feature set creation

After defining the interaction metaphor and identifying key behavioral goals, I translated the most promising design directions into feature sets. Each feature outlines a problem, articulates the user value, and defines the specific behaviors we want to support. They also clarify what data needs to be measured, what mechanisms enable these behaviors, and exactly what we would build. Together, they shape the first version of the design concept.

7.3.1. Goal Setting based on quality of life of the patient

Patients often struggle to articulate goals that align with their care plans. This feature guides them through a lightweight visual quiz that captures their current quality of life and emotional state. The system then supports them in setting Specific, Measurable, Achievable, Relevant, Time-bound (SMART) goals that feel personally meaningful.

By linking goals to lived experience, patients gain motivation and a stronger sense of control over their care. They become active participants in their treatment journey and track their progress over time. The system sends reminders, visualizes goal summaries, and prompts reflection after key treatment events. We track behavioral impact using goal completion rates, quality of life scores, and shifts in self-management attitude.

For clinicians, this feature reduces time spent clarifying patient priorities during appointments. It surfaces patient-defined goals in a clear, structured format, making shared decision making more focused and efficient. We measure success through improvements in communication quality, reduced clarification time, and stronger alignment between clinical plans and patient intentions.

7.3.2. Journal for patient mood and symptom tracking

Patients often carry emotional weight silently. This feature gives them a space to journal how they feel while tracking symptoms. It offers prompts tailored to the patient’s mood, recent events, or treatment milestones, and uses natural language processing techniques such as text analysis and sentiment analysis to detect emotional tone and shifts over time.

This journal helps patients process difficult moments and recognize patterns between emotions and physical symptoms. Over time, it supports more informed treatment discussions and healthier lifestyle adjustments. We track behavioral impact through journal engagement, stability of emotional patterns, and symptom trends. The interface allows both voice and text entries. Patients can choose to keep entries private or share summaries with their care team. This flexibility creates a low-pressure space for expression while generating valuable behavioral data that can inform personalized support.

7.3.3. Communities: Peer support matched by emotional similarity

Patients facing chronic illness often feel alone. This feature connects them with others who share similar chronic conditions and emotional patterns, allowing them to reflect, share, and receive support in a structured, anonymous, and symbolic way. This approach protects privacy while encouraging meaningful emotional connection. It helps patients build resilience, reduce feelings of isolation, and find reassurance in the lived experiences of others.

To support this, the system uses emotion-tagged prompts to spark thoughtful reflection and peer exchange. We measure success through improvements in mood trends and perceived social wellbeing. These patterns offer insight into how peer connection influences emotional health over time and generate behavioral data that can inform future support strategies.

7.3.4. Learning Hub: PAM - based information support

Uninformed or avoidant patients need structured, relevant information tailored to their readiness levels. This learning hub uses PAM questionnaires to gauge activation levels and then delivers personalized education in formats such as videos, interactive tools, and quick guides. It is imperative for the learning hub to contain clinically vetted sources of information since building trust is a significant part of the experience.

Patients receive content that fits their capability, helping them build confidence and better understand their treatment. This reduces emotional overload and supports better communication with clinicians. We evaluate progress through PAM score changes and improved comprehension. The hub includes assessments, educational paths, and follow-up quizzes that can track engagement and further act as data points for personalization.

Feature Cards

Feature set	Core problem	User Value	New Desirable behavior (What is our design space?)	Data points (How can we measure?)	Enabling Mechanism (What can we design to achieve the desired behavior?)	Action (What will we build?)
Goal setting based on quality of Life	Patients - Need assistance to set clear goals to manage their care and effectively communicate with clinicians.	1. Care feels more relevant & improves motivation. 2. QoL address emotional & social well being which is important factor missing in current care pathways 3. Patients feel in control of their care improving autonomy and ability to self-manage. This in turn improves adherence to treatment plans.	We want to help patients set SMART goals & participate in the co-creation of their care with healthcare professionals.	1. Patients actively set goals 2. Goals are actively being achieved 3. QoL scores report better health outcomes - in terms of autonomy & self-management	QoL questionnaires, Daily Goal trackers, reward systems/ motivation mechanisms	1. Lightweight visual quiz with QoL questions during goal setting (Measuring current QoL score). 2. Goal summary with daily action items(medication reminders) 3. Assess/evaluate QoL score after each intervention or treatment (Frequency varies based on disease)
Supporting Clinician dashboard	Clinicians- Need a way to understand patient behavior and goals in a structured, standardized, evidence based manner.	1. Clinicians rely not just on clinical data points but are also provided with human metrics to personalize care for chronic patients. 2. Supports shared Decision making 3. It helps make clinical work more fulfilling and patient-centered	We want clinicians to co-create and deliver care plans that are data driven and personalized based on Quality of life.	1. Reduction in clarification appointments 2. Improved communication with patients	Co-build care plans with a focus on patient goals. Implement feedback mechanisms on patients' treatment progress	1. Decision support dashboard to help doctors in building care plans that meet goals of patients (show action items) 2. Visualize adherence trends and flag deteriorating adherence for checkins 3. Provide asynchronous feedback on patient goals
Journal - for Mood and symptom tracking to take care of patient's emotional well being	Patients - Need a way to deal with uncertainty, fear and anxiety of living with a chronic disease.	1. Patients sort through events, track symptoms and reduce stress levels. 2. Puts them in a position to better deal with side-effects and later navigate treatment options. Improves and eases adjustments in lifestyle.	We want patients to take care of their emotional well being while tracking symptoms & side effects.	1. Voice/text entry frequency 2. Sentiment trends over time. 3. Logging symptoms	Adaptive journal prompts based on moods, symptoms or events. NLP based text analysis to identify sentiments	1. Voice/text entry interface with clear writing or speaking prompts. 2. A space for storing the entries. 3. Light weight Symptom/ side-effect tracker 4. Option to share insight with doctors. - not the journal, but the symptoms
Supporting Clinician dashboard	Clinicians- Need a way to understand, diagnose and treat a patients' mental well being	1. Clinicians get clarity on medication side-effects. 2. Ability to provide patients with better treatment options that are more personal. 3. A way to track mental wellbeing of their patients.	We want to promote asynchronous ways of staying involved in the care of patients to help build trust.	1. View Mood Trends over time 2. View Logged symptoms 3. Improvement in patient adherence to treatment plans 4. Patients report increased confidence in care 5. QoL report better mental well being outcomes	Flag high risk patients. acknowledge low-medium risk patients. Asynchronously adjust care plans / provide feedback for high risk patients.	1. A scannable dashboard showing mood trends and symptom summary. 2. Actionable summaries for tailoring care plans based on symptom & side-effect complaints-asynchronously. 3. A button/ mechanism to provide feedback to low-medium risk & call for consultation the high risk patients.
Communities - Peer Support for patients matched by emotional behavior and symptom similarity	Patients have to battle feelings of isolation and loneliness while maintaining emotional safety	1. Reduced feelings of isolation 2. Improves hope and resilience 3. Peer encouragement increases adherence to treatment 4. Becoming a peer supporter gives a sense of meaning and purpose	Reflecting privately on their care journey while also being socially engaged with peers/patients like them. We want to preserve the intimacy of personal journaling while allowing patients to connect and learn from others facing similar challenges.	1. Improvement in Mood trends 2. QoL report better social well being outcomes	Match patients going through similar conditions while providing a space to share & interact with peer reflections in a symbolic way (Privacy ensured)	1. A space for co-reflecting with peers while maintaining privacy. 2. Provide an option to share journal reflections with peers matched on emotional tone or topic relevance. 3. Interactions with peers/ communities (I feel this too, I have been through this, This helped me get through this..)
Learning hub - Consists of Information support modules with PAM questionnaires	Passive, uninformed and avoidant patient groups struggle with adherence to care guidelines resulting in poor health outcomes	1. Patients receive care and communication tailored to their level of engagement and capability reducing overwhelm and improving comprehension. 2. Patients feel understood and respected, which improves clinician-patient relationships	We want passive patients to take an active role, become more informed & aware of their condition, putting them in a position of making better decisions.	1. PAM Levels 2. Improved trust in care 3. Improved adherence to treatment plans	Convenient and user friendly Patient Education methods. PAM questionnaires, PAM level badges.	1. PAM assessment questionnaire for patients 2. An information support page to improve patient education designed to be inclusive of all patients 3. Options and links to additional information support
Supporting Clinician dashboard	Passive, uninformed and avoidant patient groups are seen across the three diseases. Clinicians need a way to identify & keep track of this important patient group.	1. Clinicians get insights into a patient's ability to self-manage their health, enabling more personalized treatment plans that align with PAM levels. 2. Time used in consultation room is better utilized by focusing on what the patients need.	We want clinicians to Identify & keep track of low activated patients and tailor communication and support strategies in a way that empowers the patient and encourages self-management.	1. PAM levels 2. Time in consultation rooms(reduced)	Provide/direct patients towards Information support including (FAQs, reading links, vetted videos, Quiz of the day + information)	1. Dashboard indication current PAM levels(improved/stagnant PAM indicators) 2. Provide materials to improve patient education at the click of a button/brochure etc 3. Mark or filter patients by activation levels & risk

Fig 7.4. Feature sets table

7.4 Designing the user flows for the Patient

After finalizing the feature cards, I translated user goals and actions into a structured flow that guides the design of the patient-facing application. The user flow diagram (see fig 7.5) maps how patients interact with each feature, move through key decision points, and perform meaningful actions. The diagram reflects every major function outlined in Section 7.3 and provides a blueprint for the design concept.

I designed this experience with clarity and ease of use in mind. The flow aligns with the interaction principles defined earlier supporting progressive onboarding, maintaining emotional warmth, and ensuring that every step leads to value.

To make the user journey easier to interpret, I created a consistent visual structure. Each screen begins with a labelled screen name that describes its purpose. Inside each screen, I divided the interface into page sections for easier mapping. Under each section, I added components, which describe how specific elements function within the screen.

When a user interacts with a feature, the flow shows it as an action card. If the system responds in the background, such as analyzing quiz responses or uploading a symptom entry, this appears as a data action. Finally, any point where the user makes a decision appears as a decision node, represented by a diamond.

This framework allowed me to track not only where users go but also what they think and choose at each step. The next pages walk through these flows feature by feature, beginning with the onboarding experience for first-time users.

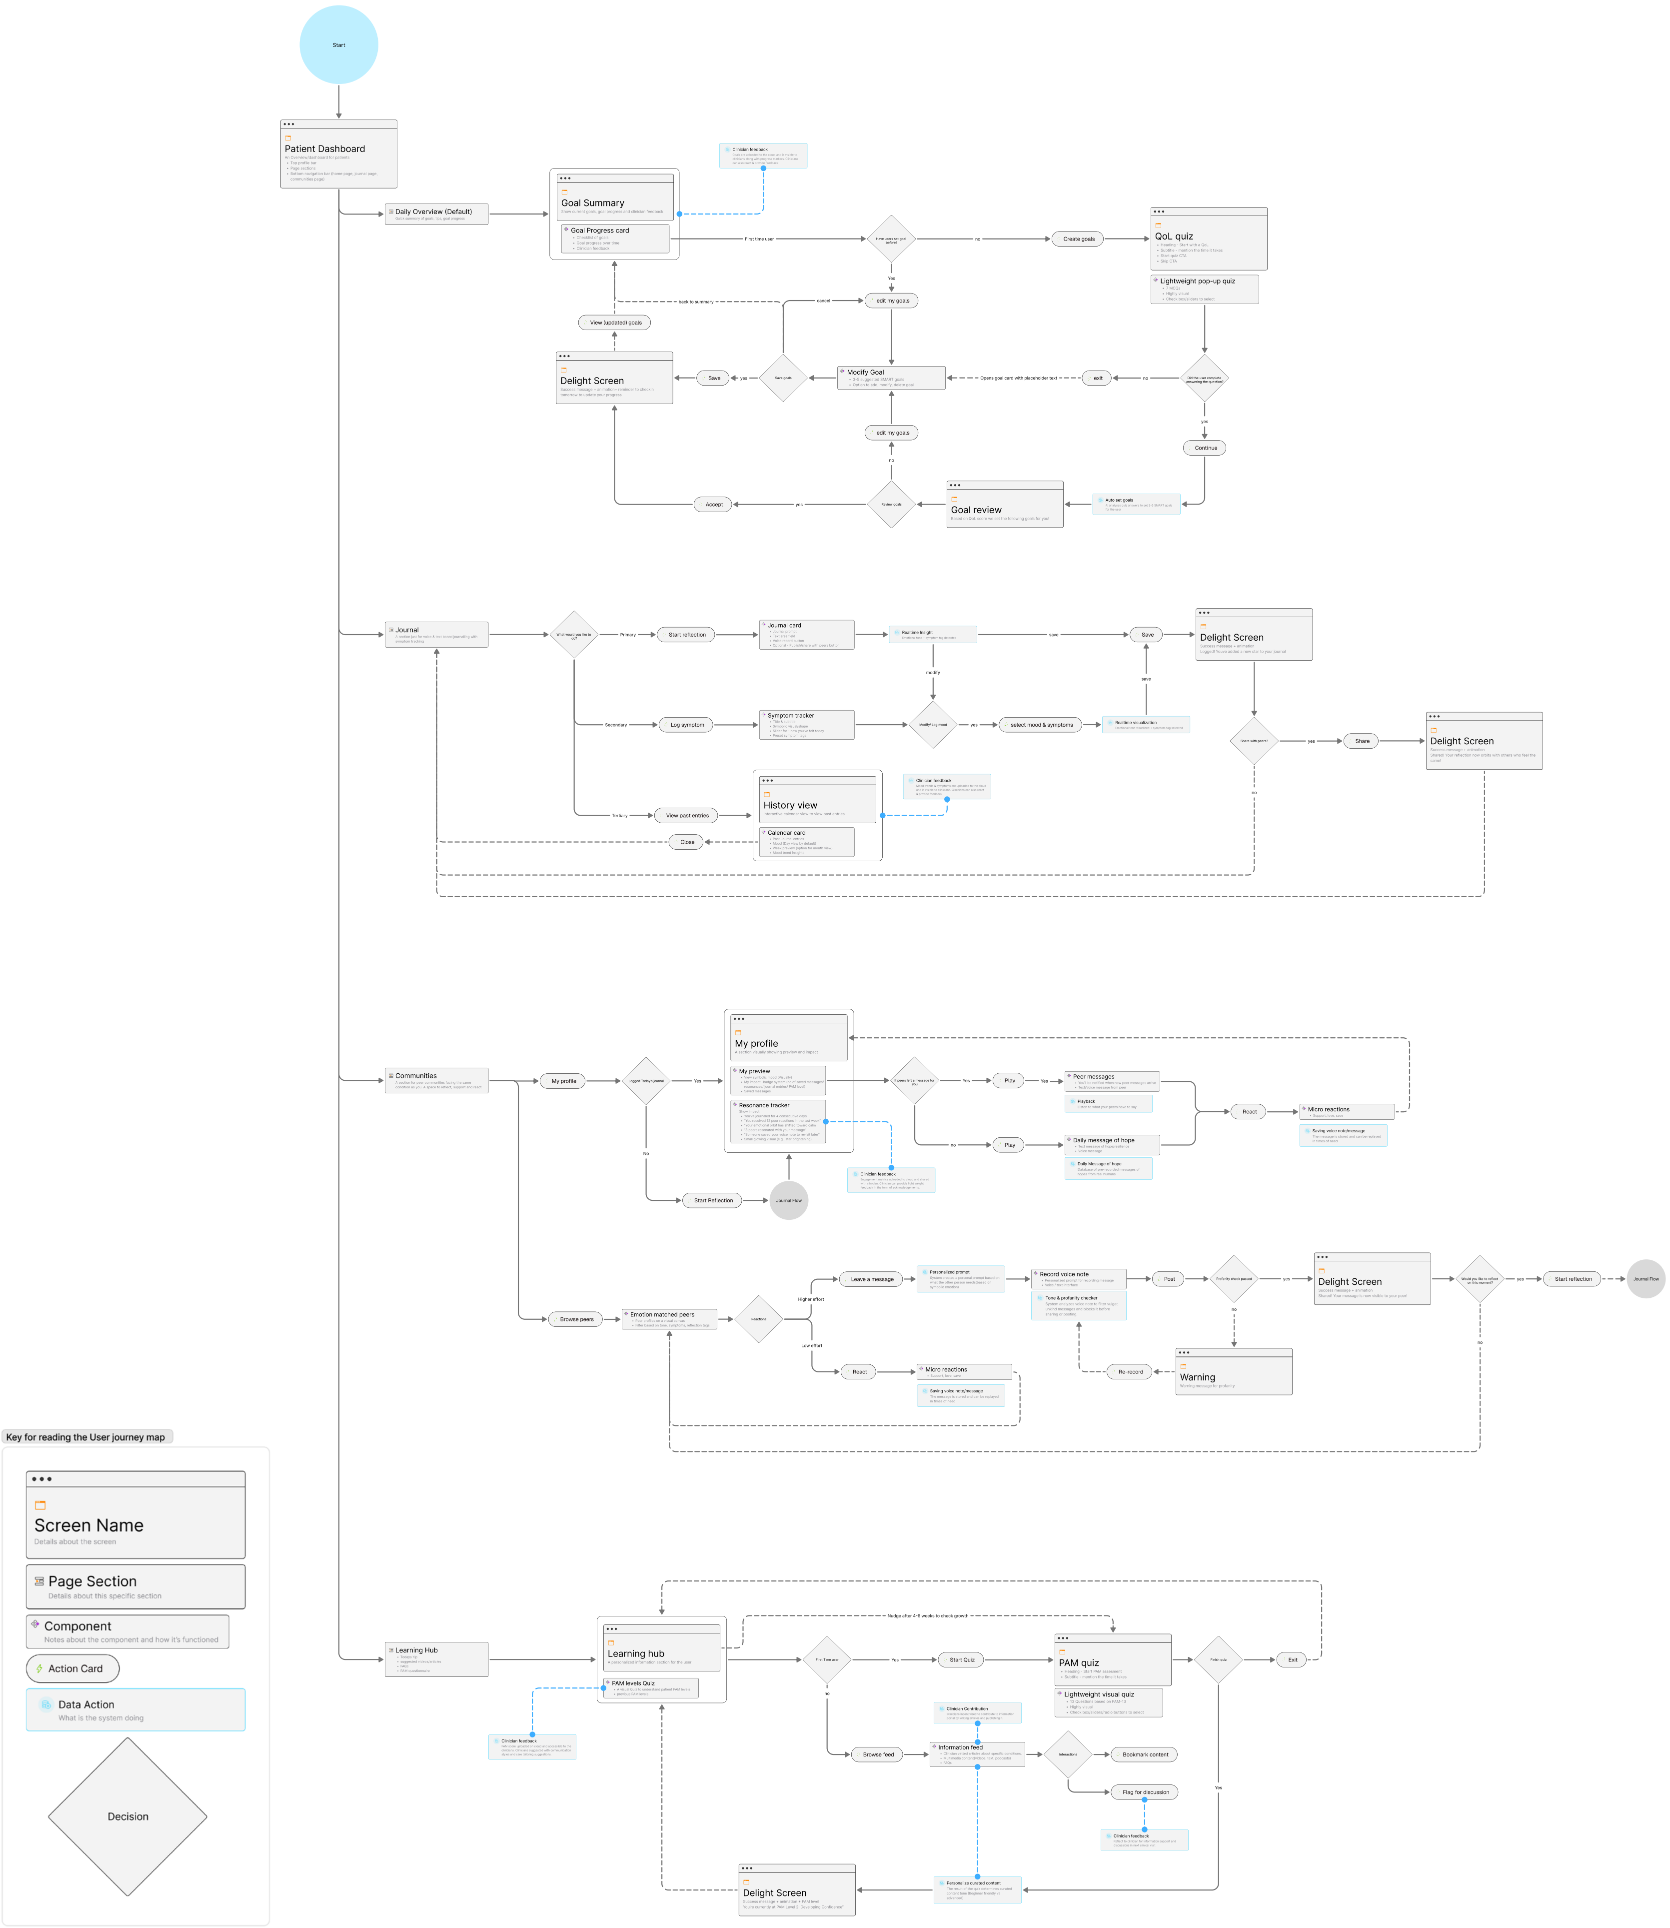

Fig. 7.5 Patient facing app user flow diagram

7.4.1 Onboarding Flow (First Time User flow)

The onboarding flow introduces new users to the app and helps them set up their basic profile. It begins with a welcome screen that displays a short message, a subtitle, and a calming visual illustration to create a warm first impression.

Next, users reach the account setup screen, where they sign up using a mobile number or email address. This screen includes components for entering login details, verifying via One Time Password (OTP) or password, and agreeing to terms and conditions. At this stage, users also provide consent for personal data processing, a key requirement that aligns with GDPR and the EU AI Act (Artificial Intelligence in Healthcare, 2025). After account creation, users move to the condition select screen. This page allows them to identify their current condition using a visual selection grid. They can either choose a specific diagnosis or select “I don’t know my condition yet” if they are still awaiting confirmation.

Once this step is complete, a delight screen marks the end of onboarding with an encouraging message and animated illustration. This small moment of celebration gives users a sense of completion and invites them to begin exploring the app.

The flow ends with the patient dashboard, which acts as the home screen. It provides a summary of goals, progress, and navigation to the app’s main features. Key user actions in this flow include creating an account, selecting a condition, and giving consent for data use. The central decision point occurs during account setup, where users must choose whether to proceed with data sharing.

Designing the flow in this way brings three main benefits:

1. It reduces friction by asking only for essential information
2. It personalizes the user experience early without creating overwhelm
3. It reinforces a sense of emotional support and clarity from the very first interaction

This onboarding flow sets the foundation for a personalized and emotionally supportive user experience from the very first interaction.

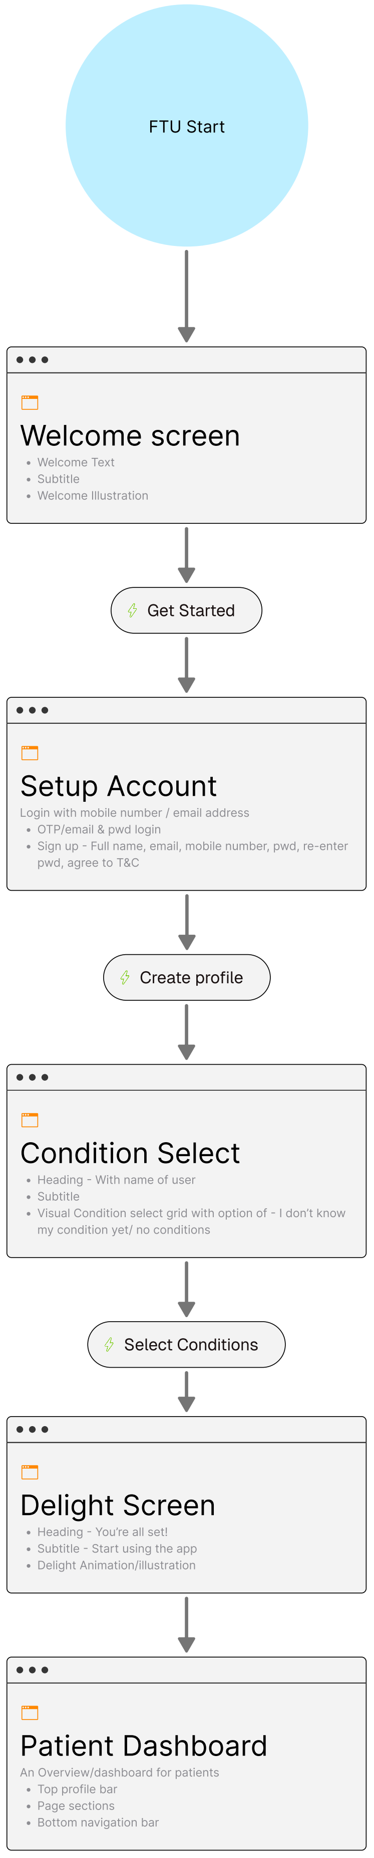

Fig 7.6 First Time User Flow & onboarding

7.4.2 Goal Setter Flow

The goal setter flow helps patients define meaningful goals based on their current quality of life. It starts with a Quality of Life Quiz, where users respond to a short series of visual questions using sliders and checkboxes. This quiz collects behavioral context and emotional insights to inform the goal-setting process. The quiz will be based on an already existing questionnaire on patient reported outcomes such as the short form health survey or the WHOQOL- BREF (WHOQOL-BREF, n.d.; 36-Item Short Form Survey Instrument (SF-36), n.d.).

These questionnaires are condition specific and appropriate clinical guidance is necessary for selecting these questionnaires. The answer to these questions will be used as an input parameter for the backend systems to act as data points for the creation of SMART goals (Specific, Measurable, Achievable, Relevant and Time-bound)

After submitting their responses, users reach the goal review screen, where the system uses the input parameters of the user to suggest 3 to 5 SMART goals tailored to their quiz results. Users can edit, add, or delete these goals through a simple interface that allows flexible customization. This supports autonomy and ensures the goals reflect what each patient truly cares about.

Once goals are confirmed, a goal summary screen shows an overview of current goals, progress indicators, and space for clinician feedback. This screen encourages users to keep track of their progress and understand how their actions align with care plans. Users then see a delight screen, which reinforces their progress with a success message and a reminder to check in again the next day. This positive feedback loop helps sustain motivation and builds confidence in self-management. The flow ends on the daily overview screen, where users can return each day to review goal progress, receive personalized tips, and stay on track.

Key user actions in this flow include completing the quiz, reviewing suggested goals, and confirming or modifying them. Decision points occur when users choose to accept, edit, or skip goal suggestions. This branching moment helps personalize the experience based on user readiness and engagement level.

Designing the flow this way ensures patients feel seen and supported. The system builds trust by setting achievable goals, framing progress visually, and reinforcing motivation through small moments of success. The ability for clinicians to provide their inputs on the goals set by patients adds a layer of trust and acknowledgement. This flow empowers patients to set goals that align with their emotional and physical needs, making care feel actionable, achievable, and personal.

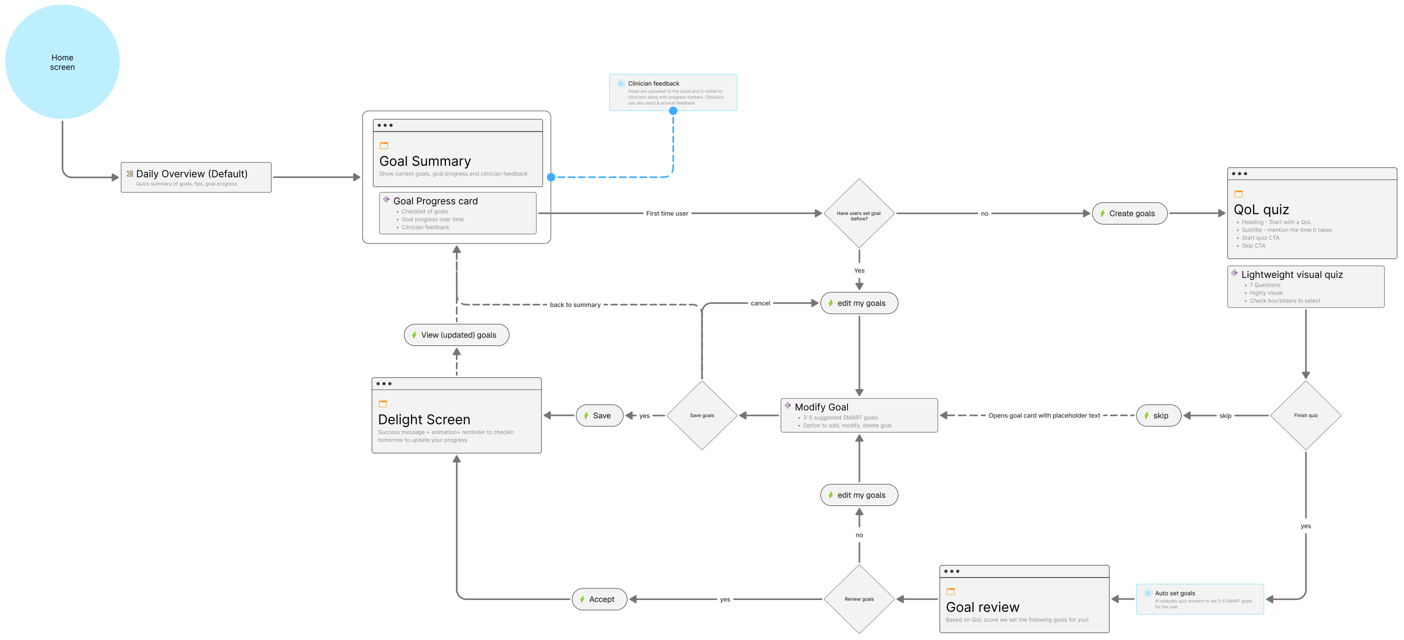

Fig 7.7 Goal Setter flow

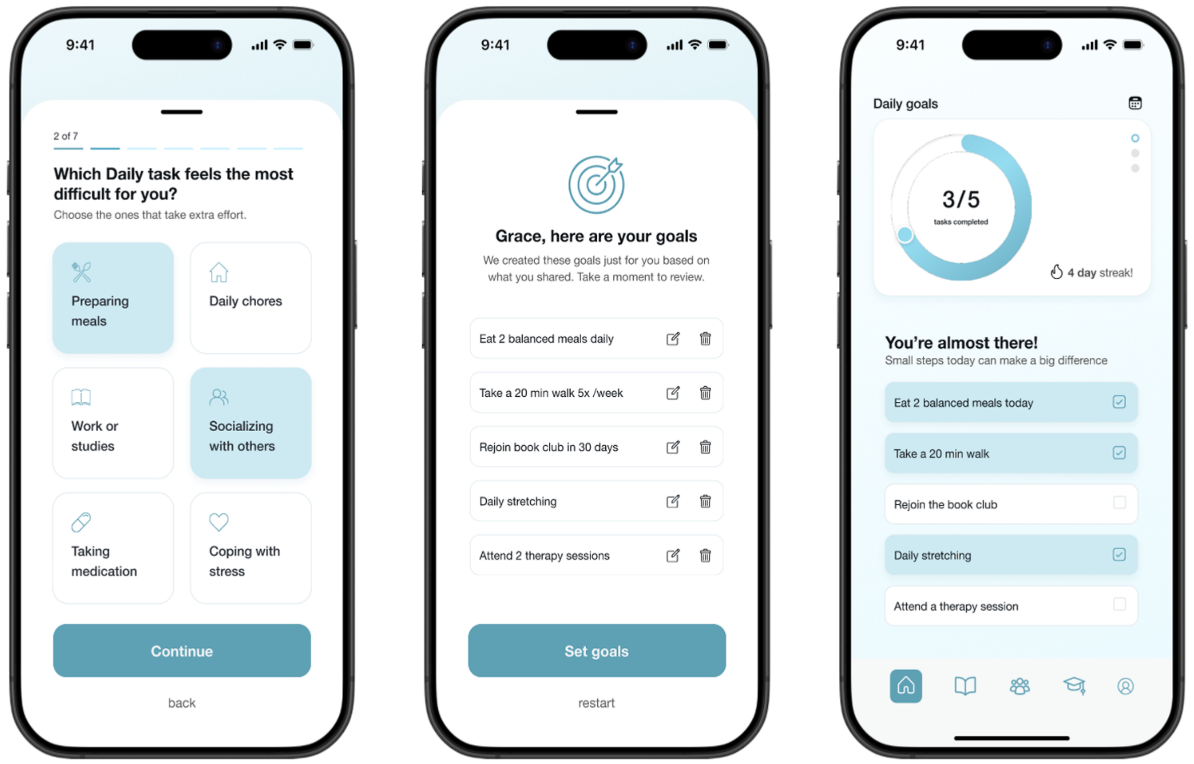

Fig 7.8 Goal Setter flow key screens; from left to right (Goal setting questionnaire, Review goals, Daily tasks)

7.4.3 Journal Flow

The journal flow gives patients a space to reflect on their emotions and track symptoms in a low-effort, supportive format. It starts with the journal entry screen, where users see a prompt and choose whether to write a reflection, record a voice note, or log symptoms. The screen includes a mood slider, symbolic visuals, and preset symptom tags that help patients express their experience clearly and quickly.

Users can choose to save the entry privately or share it with peers through the publish option. If they decide to share, the system checks for emotional tone and filters for harmful or insensitive content using a built-in sentiment and profanity filter. After the entry is saved or shared, the user sees a delight screen that affirms the action. For private entries, the message reads "Logged! You've added a new star to your journal." For shared entries, it reads "Shared! Your reflection now orbits with others who feel the same."

The flow also accounts for the history view, which displays past entries in an interactive calendar format. Users can track their voice journals, mood and symptoms over time and access these insights through day, week, or month views. This retrospective view supports self-awareness and helps users recognize patterns.

Key user actions include starting a reflection, selecting mood and symptoms, deciding whether to share, and reviewing past entries. The primary decision point occurs when users choose to log privately or publish to the peer community. This moment offers both introspection and social connection.

Designing the journal flow in this way supports emotional safety and behavioral insight. It gives patients a consistent and judgment-free space to express how they feel. The system analyses emotional tone in real time and stores mood and symptom trends in the cloud, where clinicians can view them through their dashboard.

This flow enables users to externalize emotion, observe their own patterns, and connect with others when they choose to. It builds emotional resilience through reflection and creates valuable behavioral insights for adaptive care.

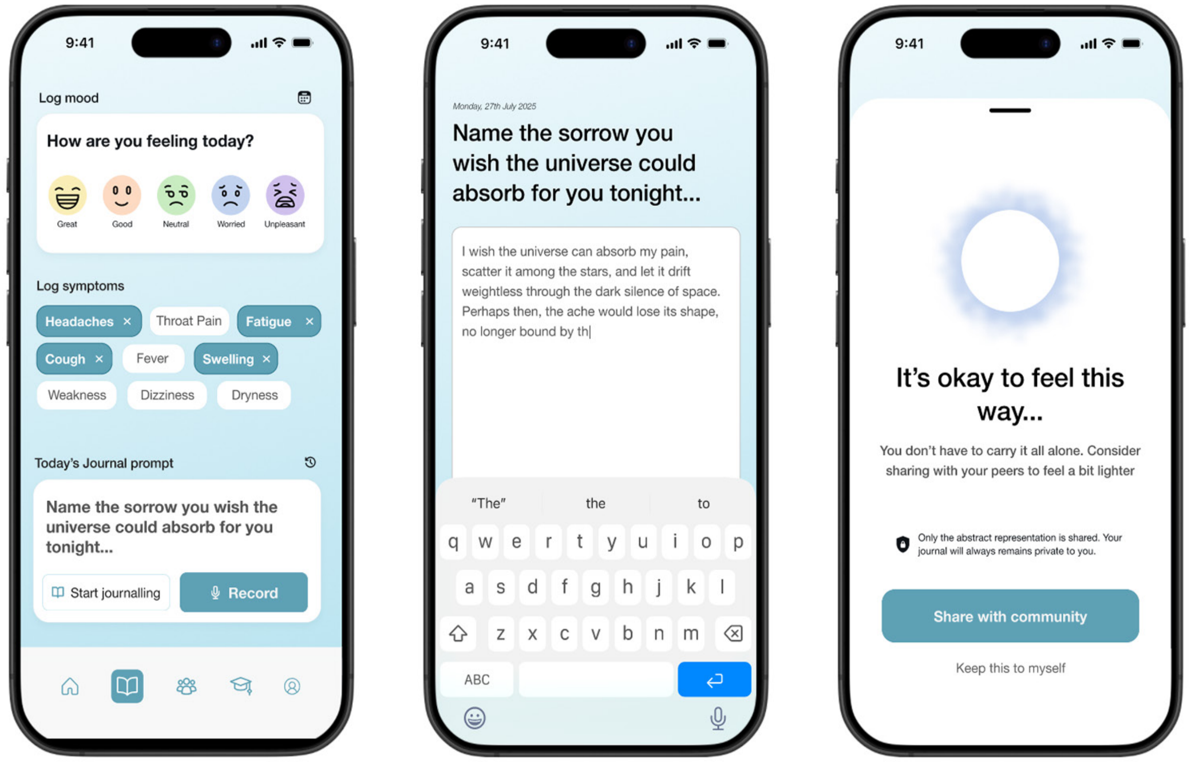

Fig 7.10 Journal flow key screens; from left to right (Lightweight mood logging, Journalling, Abstract representation)

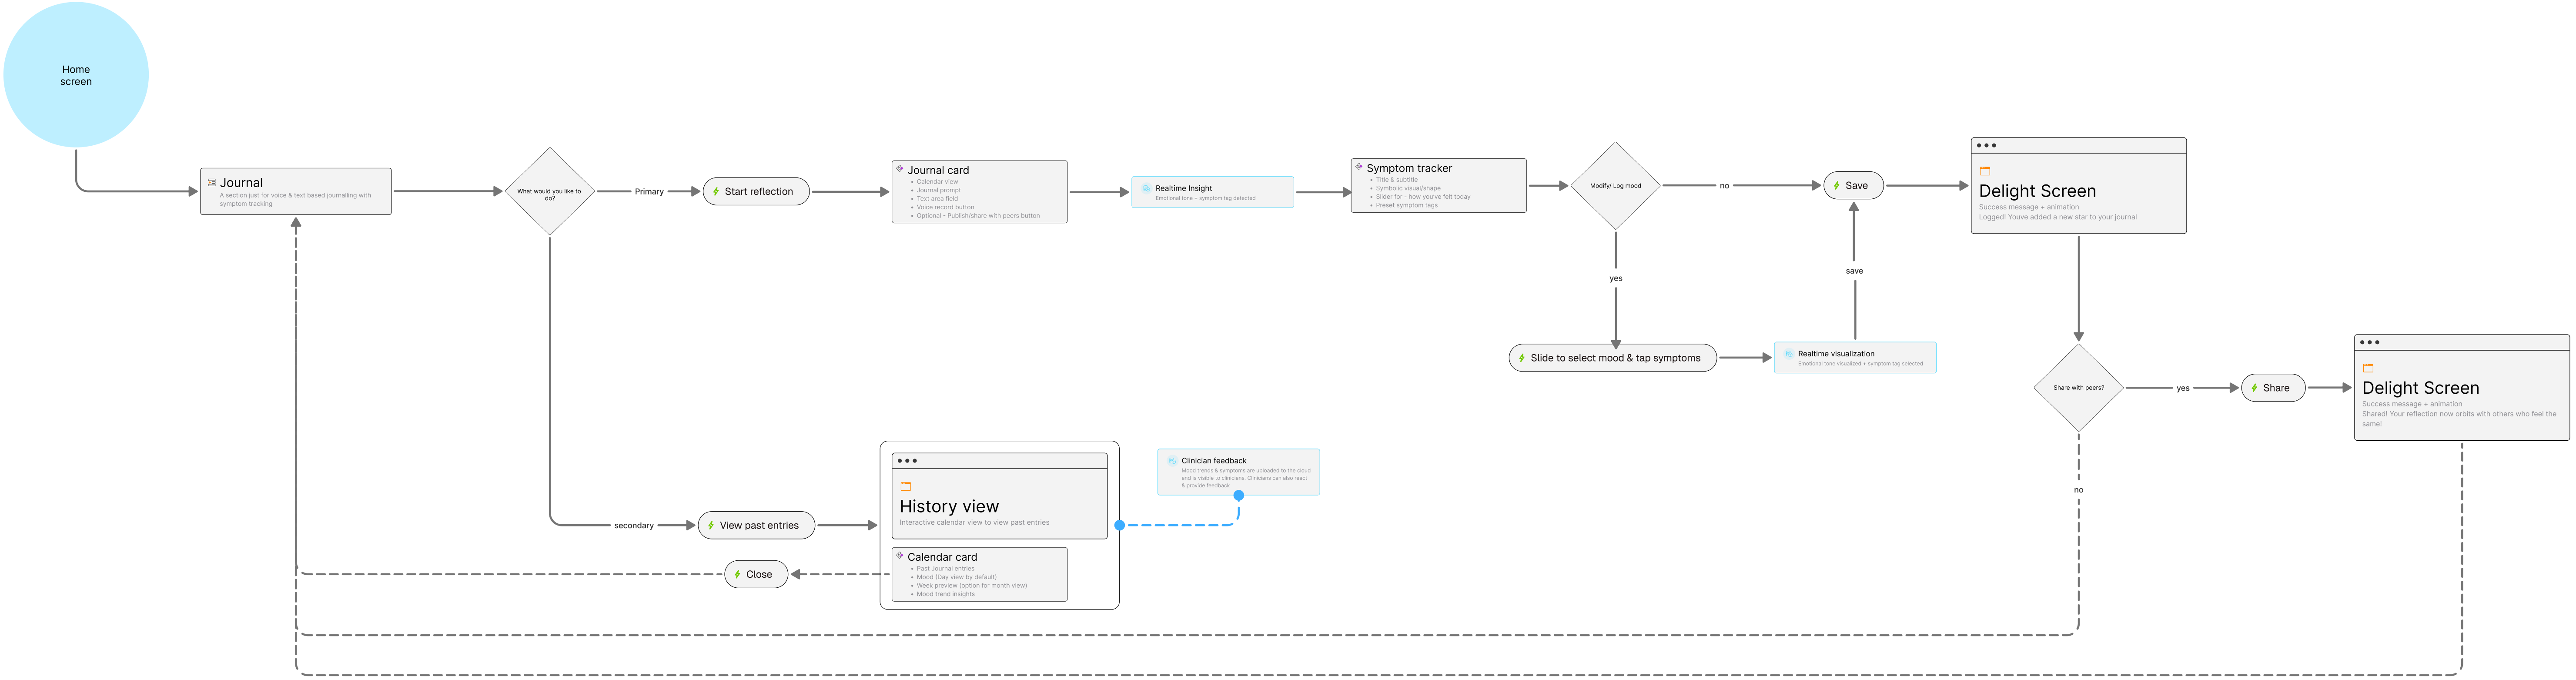

Fig 7.9 Journal flow

7.4.4 Community Flow

The community flow introduces a novel and emotionally intelligent space within the app. It explores the creative tension between holding an experience privately and choosing to share it symbolically with others. Designed as a bridge between solitude and connection, this feature invites patients to engage in peer support without forcing exposure or overstimulation.

The experience begins in the communities section, where patients enter a calm, reflective environment curated around shared emotional states rather than fixed diagnoses. Here, they can explore Emotion-matched peer profiles, visualized as symbolic representations on a canvas that resembles a constellation. These visual markers reflect mood, tone, and thematic similarity, allowing patients to connect with others who feel emotionally familiar, even if their conditions differ.

Patients who wish to contribute can leave reflections through a voice or text Interface. The system generates a personalized prompt based on the user's current mood or symbolic state. This step is low pressure and deeply flexible. Before sharing, the system activates a tone and profanity checker to ensure safety and emotional integrity. Once cleared, the user sees a delight screen confirming, "Shared! Your message is now visible to your peers."

This act of sharing transforms the private into the collective without stripping away personal meaning. A journal entry that once felt inward now becomes a point of resonance. Peer reflections arrive in return as daily messages of hope or spontaneous notes of encouragement. These may include voice recordings or short texts, curated based on emotional relevance rather than random matches.

Users can respond with micro reactions (support, love, or save) keeping interaction lightweight but emotionally significant. All interactions feed into a resonance tracker visible symbolically. Patients receive symbolic feedback in the form of a glowing star, a rising badge count, or a note that someone saved their message to revisit later. These subtle reinforcements build a sense of emotional contribution and mutual care.

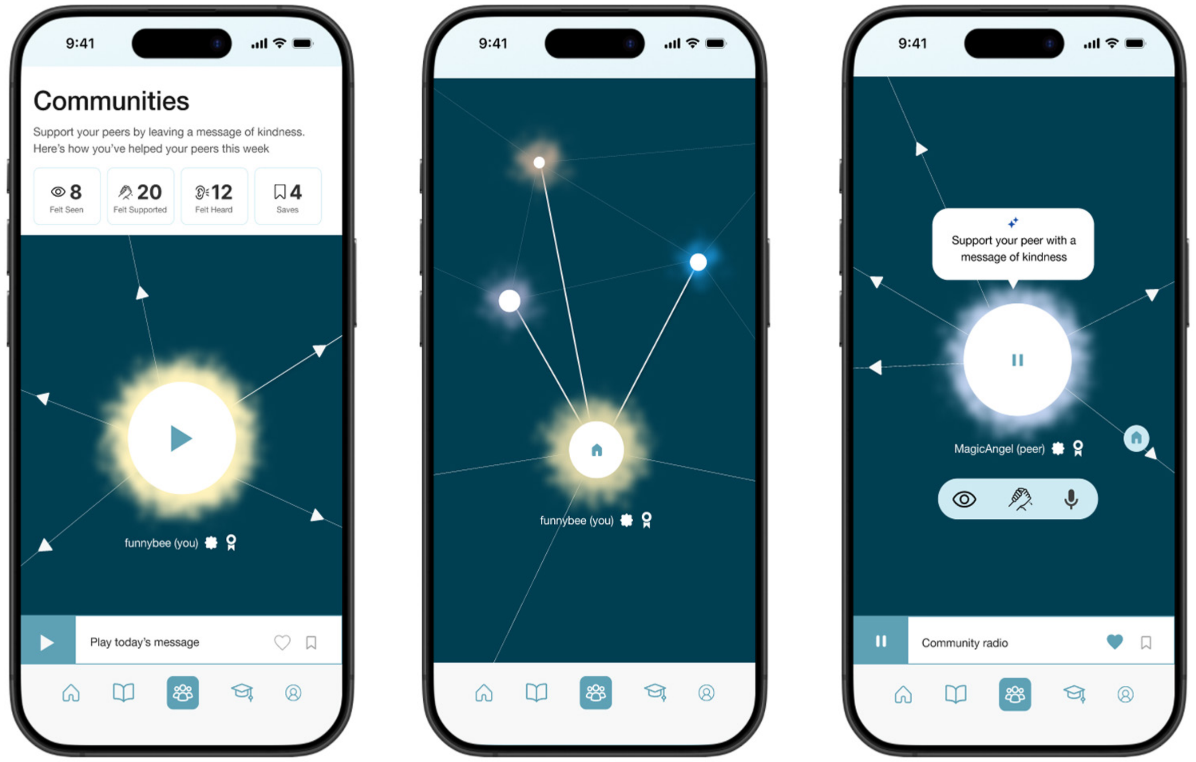

Fig 7.12 Community Flow key screens; from left to right (Community home, Explore peers, Support and interact with peers)

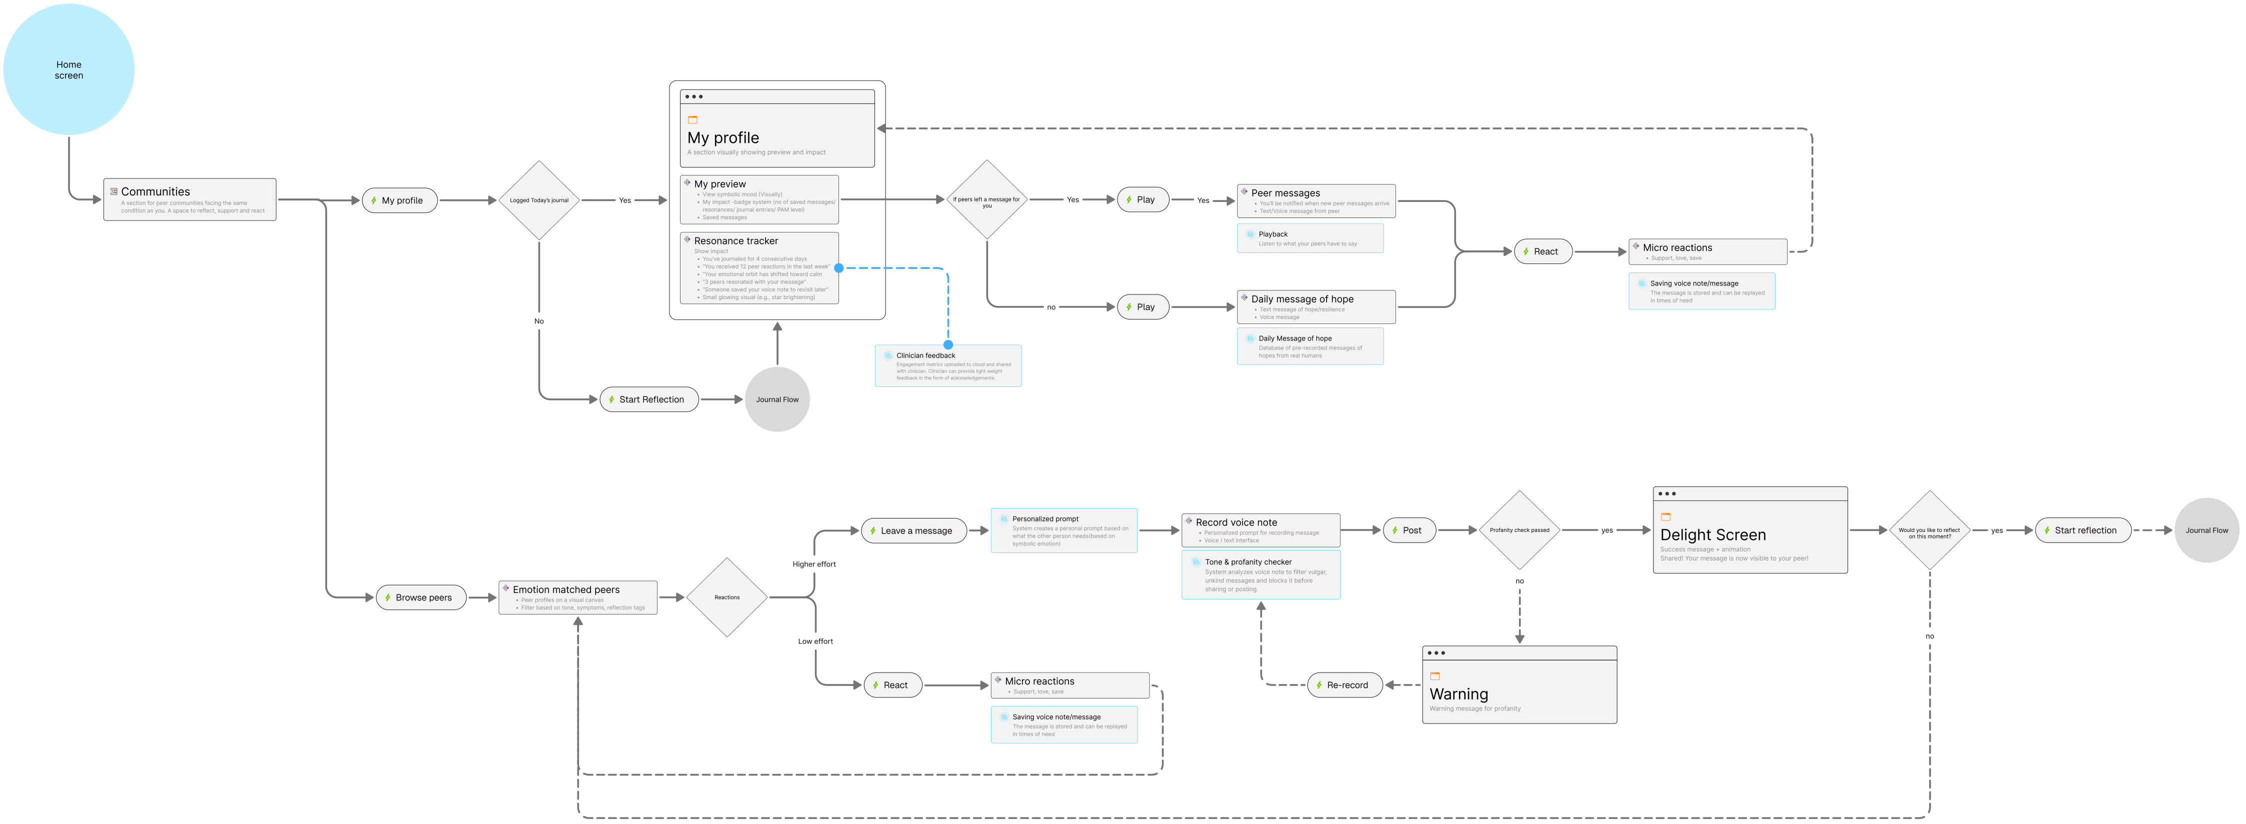

Fig 7.11 Community Flow

Key user actions in this flow include browsing symbolic peers, choosing to share reflections, responding to messages, and viewing one's emotional impact over time. A key decision point arises when the app asks whether to share a journal entry with peers or keep it private. This moment balances vulnerability with agency, allowing patients to decide how much of their inner world they want to place in front of others.

This flow redefines peer support through emotional symbolism rather than direct messaging. It encourages reflection over reaction, recognition over judgment, and resonance over instruction. It does not attempt to replace therapy or clinical dialogue however, it creates a quiet, human space between them.

This flow introduces a new design language for peer connection. By using symbolic emotion and shared reflection as the medium, it gives patients a safe way to feel seen and valued. In doing so, it transforms private expression into a shared constellation of care.

7.4.5 Learning Hub Flow

The learning hub flow gives patients access to personalized health information that matches their current level of readiness and confidence. It focuses on delivering clear and relevant content instead of overwhelming users with generic advice.

The flow begins with a PAM Level Quiz, a short visual questionnaire that helps understand how confident the patient feels in managing their condition. The quiz will be based on the PAM-13 questionnaire that is already widely used in clinical research (Roesel et al., 2024). The quiz is designed to use checkboxes and sliders, in a highly visual format to make it easy to complete. After finishing the quiz, patients arrive at a delight Screen while being placed in a PAM Level from 1 to 4. A delightful message like “You are currently at PAM Level 2: Developing Confidence” acknowledges their progress and sets a positive tone.

Next, the patient enters the Learning Hub, which presents a curated feed of content. The app uses the PAM score to adjust what the user sees. Patients with lower scores receive beginner-friendly content such as short videos, infographics, and simple FAQs. Patients with higher scores access in-depth articles, podcasts, and more complex materials. This approach supports learning that feels relevant and manageable without overwhelming the user.

Patients can take actions such as bookmarking content, saving items for later, or flagging topics to discuss with their clinician. After a few weeks, the app sends a gentle prompt to retake the quiz and track personal growth. As patients engage with the learning materials, the system adjusts the feed to match their evolving needs.

Clinicians receive the PAM scores and engagement metrics on their dashboard. This gives them insights into how well the patient understands their condition. Clinicians can offer suggestions, adjust communication styles, or update care plans based on what the patient has read or flagged for discussion. The flow also invites clinician contribution. Clinicians can create and publish educational content that the system recommends to specific patients. This keeps the learning experience rooted in trusted sources while allowing it to stay personal and adaptive.

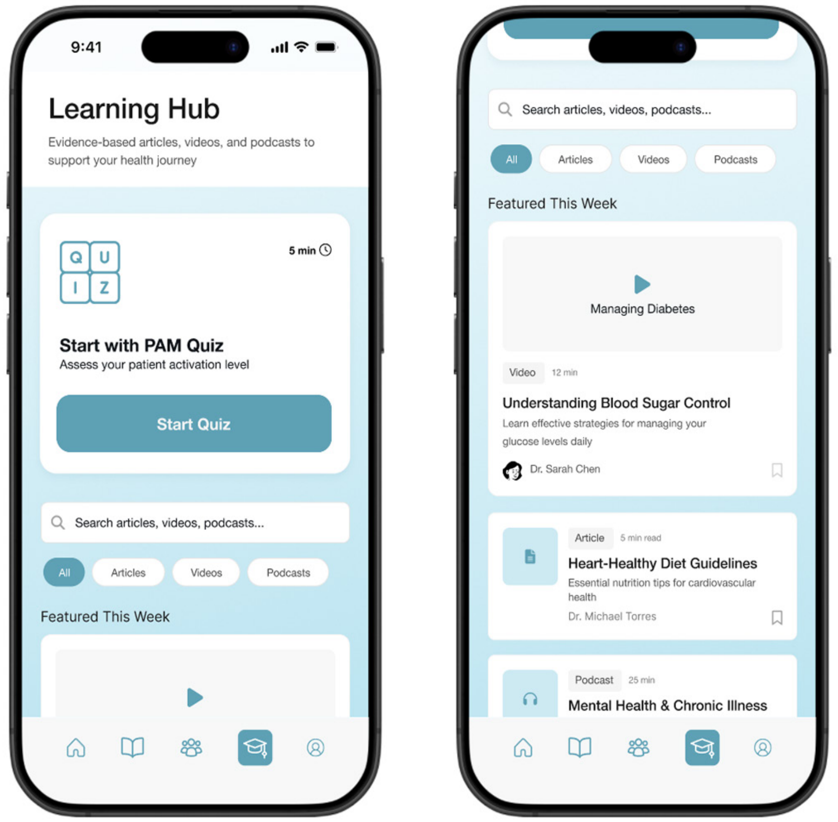

Fig 7.14 Learning hub key screens; from left to right (PAM quiz & search functions, multi-media content)

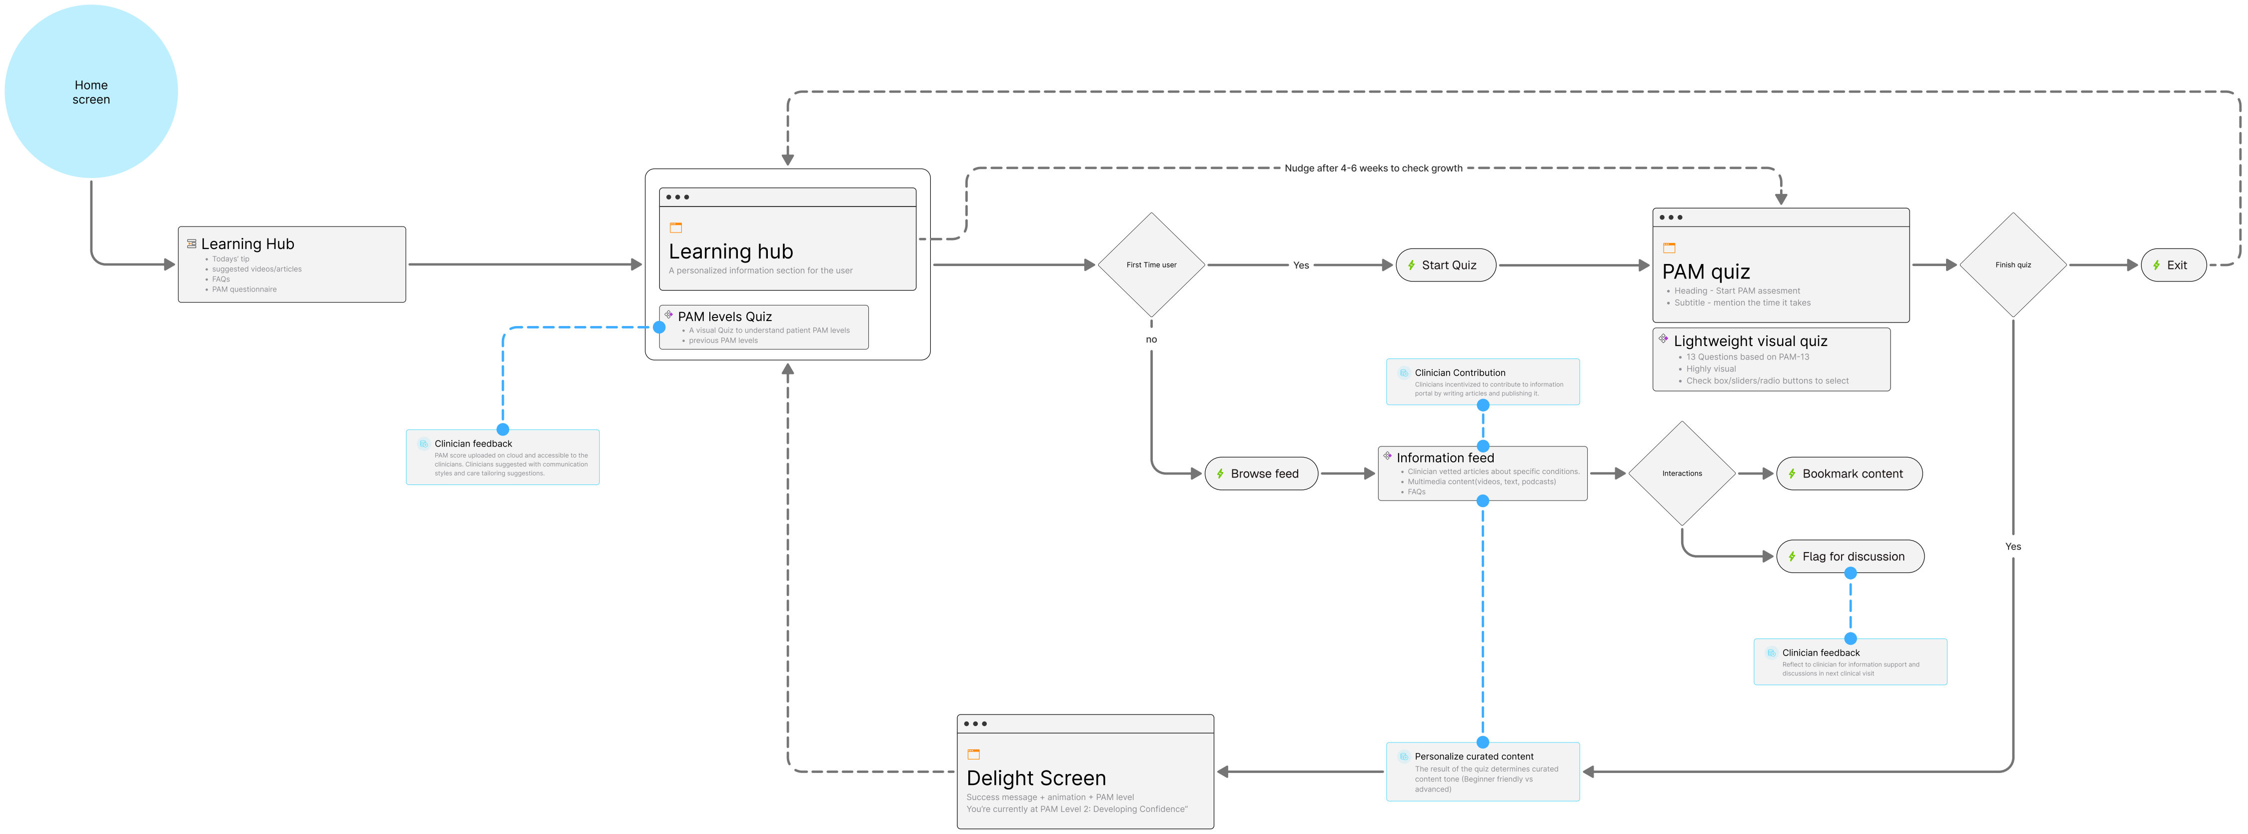

Fig 7.13 Learning hub Flow

Key user actions include completing the quiz, exploring content, saving materials, and marking topics for clinical discussion. The main decision point happens when the user chooses to begin the quiz. This choice starts the process of personalization and builds the foundation for future learning.

This flow helps patients grow their understanding over time. It delivers information when the patient feels ready and keeps the learning journey gentle, focused, and easy to follow. The learning hub flow supports each patient’s progress by giving them the right information at the right moment. It encourages continuous learning, improves patient confidence, and strengthens the connection between education and care.

7.5 Reflections

This phase of my project challenged me in ways I did not expect. I knew the direction I wanted to move in, but translating systemic insights into specific features felt like walking through fog. I kept revisiting the same questions and second-guessing each decision. At one point, I felt completely stuck. The connection between patient behavior, clinical workflows, and interface design felt distant, even though I had spent months immersed in each of them.

I reached a point where pushing harder did not help. So, I paused. I stepped away from my screen and spent time in nature. I walked without an agenda. Being outside helped clear the mental noise and reset the pace of my thinking. It reminded me that clarity often comes when I stop trying to force it. I took a notebook with me and started capturing fragments of words, images, and small feature ideas. Later, I got on a treadmill with a friend and talked it out. That conversation gave me the kind of feedback I could not get from staring at my work. It felt informal and honest. I spoke freely, and in doing so, I unlocked parts of the problem I had not seen before. After that, I returned to the whiteboard. I created rough flows and shared them quickly with a few peers for quick feedback. This fast feedback loop helped me shape the details without overthinking them.

Most of all, I reminded myself to trust the process. Design does not move in a straight line. It unfolds through rhythm, through tension, pause, and reconnection. Letting go of the need to perfect each step gave me space to return with sharper focus. This chapter taught me to design with presence, not pressure. To listen to my own instincts more closely and to stay rooted in the real human problems that sparked this project in the first place.

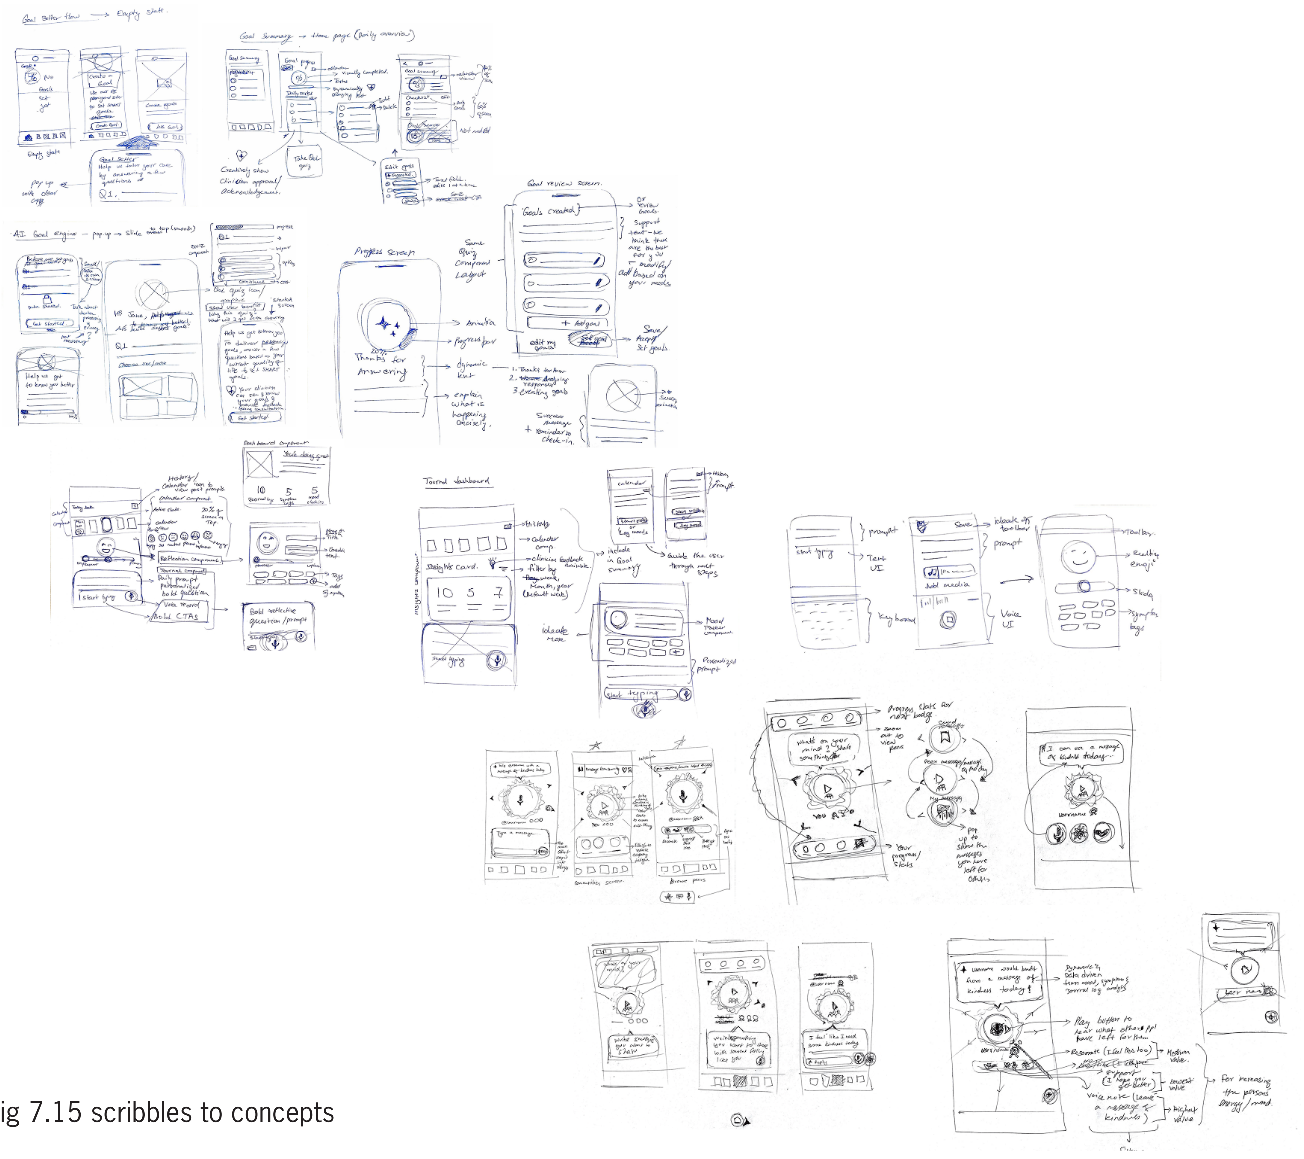

Fig 7.15 scribbles to concepts

7.6 Conclusions

This chapter explored how we might begin to design patient-facing tools that capture behavioral data in ways that feel supportive and personal. I translated emotional pain points into a set of digital features that help patients reflect, set goals, connect with others, and learn at their own pace.

The first research question asked how we could turn patient pain points into useful features. I responded by designing tools that speak to what patients actually go through—uncertainty, isolation, and emotional stress. Each feature grew from those lived experiences and tried to offer something clear, manageable, and meaningful in return.

The second question looked at how we might collect behavioral data without adding pressure. I focused on building interactions that feel natural and low effort. The system prompts patients gently, respects their choices, and lets them decide what to share. At the same time, it captures valuable insights that can support more personalized care.

The third question asked how personalized flows can support engagement. I used quality of life scores and activation levels to shape what patients see and do. These flows help patients feel seen and understood. They also build confidence in patients by making small wins visible.

The final question focused on how this data might help clinicians. While the design and development of the clinician facing dashboard comes in the next chapter, this design work sets the stage. It creates structured outputs like mood trends, goal progress, and learning activity insights that can help clinicians respond with more empathy and precision.

This chapter was not about building a finished product. It was about imagining what this future might feel like for the patient. It helped me move from broad vision to grounded interactions. Each feature became a small piece of that future, shaped by human factors such as emotion, behavior, and trust.

In the next chapter, I turn to the clinician side of the story. I design and test a future-facing dashboard that brings these insights into clinical conversations.

Chapter 8 Clinician Dashboard Design

The previous chapter explored how patient-facing tools can collect behavioral and emotional data to support self-management. But collecting data is not enough. To create meaningful impact, clinicians need intelligent tools that help them interpret this information and apply it during care planning. Without this, patient insights remain disconnected from clinical decisions.

This chapter focuses on designing that missing clinical layer. It introduces a prototype of a future-facing dashboard that functions as an agent within the clinician's workflow. The interface turns patient-generated data into care modules that clinicians can review, adjust, and assign. It also helps them communicate care plans in language that patients can understand and trust. The goal of this chapter is to explore how agentic interfaces might support personalization in ways that reduce clinical burden. I investigate three key questions

"How can behavioral data be visualized to support clinical decision making?"

"How might we reduce the burden of personalizing care from a clinician's perspective?"

"How can AI support personalization while preserving clinician autonomy?"

Through this work, I bring the central research question of the thesis into sharper focus. I explore how data-enabled design can support the creation of personalized care pathways that respond to the lived experiences of chronic patients. At the same time, I look at how intelligent tools can reduce the cognitive and operational load on healthcare providers. The agent interface serves as a test case for what this future might look like in practice.

This chapter also describes my prototyping workflow using Cursor AI. I explain how I defined the feature set, built interface components, and refined the design through iteration. The chapter ends with personal reflections and sets the stage for expert validation in the next phase.

8.1 Clinician Dashboard UX

In previous projects involving the design of digital health tools, clinicians consistently highlighted the need for solutions that integrate into existing hospital systems such as the HIX platform (Quijada Fernández, 2024). Based on this feedback, I designed the clinician dashboard as a plugin that fits directly within HIX, which remains the most widely used electronic health record system in Dutch hospitals. The dashboard follows a single-screen structure that emphasizes clarity, relevance, and speed. It applies universal principles of interface design to minimize cognitive load and increase efficiency (Lidwell et al., n.d.).

Clinicians work under time pressure and need interfaces that give them the right information at the right moment. The dashboard uses a fixed three-zone layout. The left sidebar displays patient identity and appointment navigation. The center panel shows the main clinical content, beginning with the Clinical Summary tab. The right panel supports notetaking and documentation with an AI assistant.

When clinicians open the dashboard, they first see the Clinical Summary. This view presents essential information including diagnosis, staging, allergies, vitals, and medications. The care pathway generation button appears at the top right of this view. This allows clinicians to generate AI-supported care modules based on current clinical and behavioral data.

This section of the dashboard keeps the focus on medical history and physiological data. It grounds behavioral insights in clinical context. PAM levels and patient identity appear in the sidebar to help clinicians quickly recognize the patient's self-management confidence.

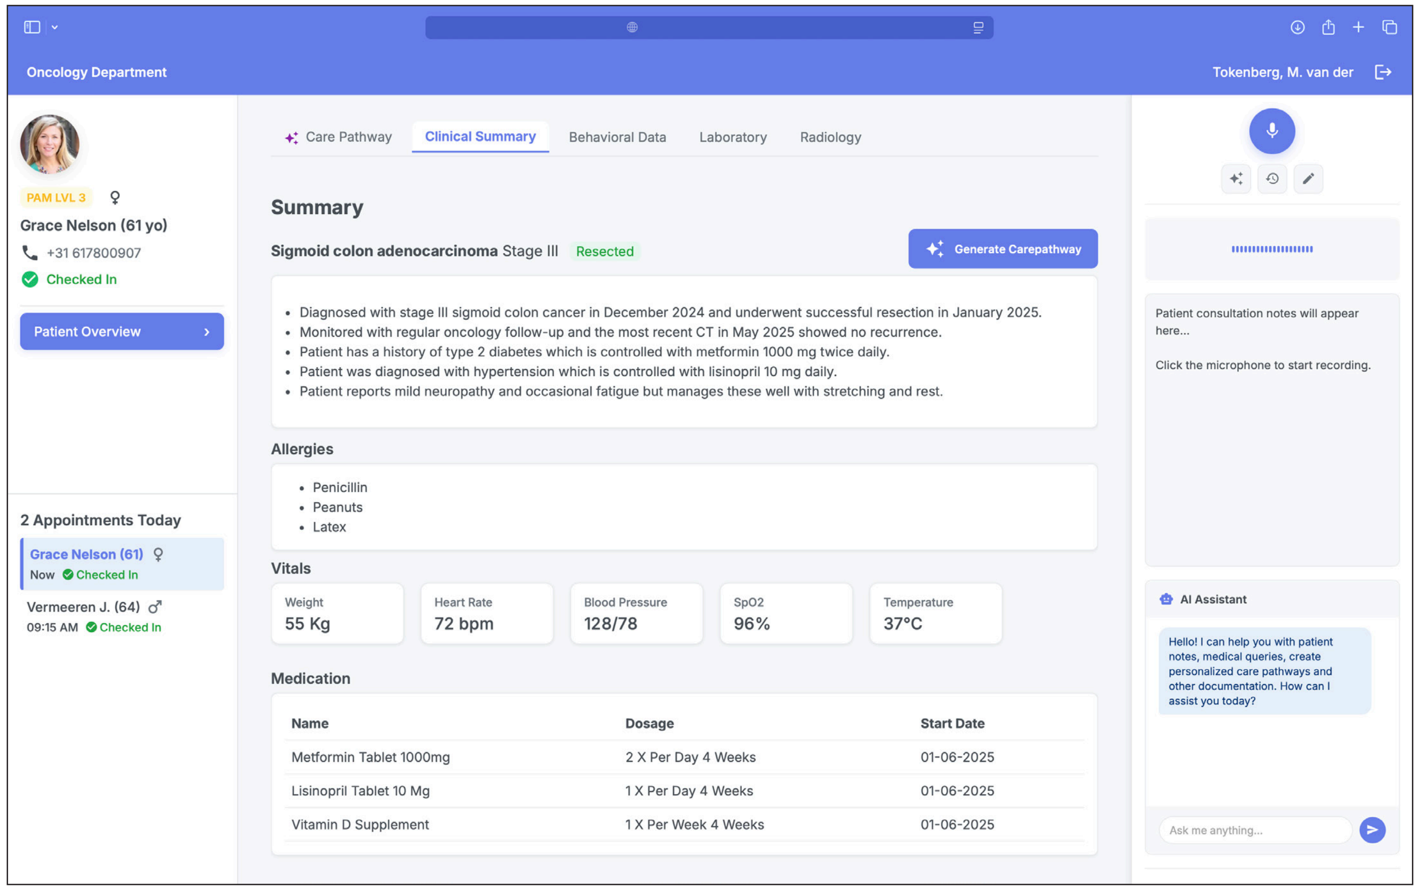

Grace Nelson (61 yo)
+31 617800907
Checked In

2 Appointments Today
Grace Nelson (61)
Now Checked In
Vermeeren J. (64)
09:15 AM Checked In

Summary

Sigmoid colon adenocarcinoma Stage III Resected

- Diagnosed with stage III sigmoid colon cancer in December 2024 and underwent successful resection in January 2025.
- Monitored with regular oncology follow-up and the most recent CT in May 2025 showed no recurrence.
- Patient has a history of type 2 diabetes which is controlled with metformin 1000 mg twice daily.
- Patient was diagnosed with hypertension which is controlled with lisinopril 10 mg daily.
- Patient reports mild neuropathy and occasional fatigue but manages these well with stretching and rest.

Allergies

- Penicillin
- Peanuts
- Latex

Vitals

Weight	Heart Rate	Blood Pressure	SpO2	Temperature
55 Kg	72 bpm	128/78	96%	37°C

Medication

Name	Dosage	Start Date
Metformin Tablet 1000mg	2 X Per Day 4 Weeks	01-06-2025
Lisinopril Tablet 10 Mg	1 X Per Day 4 Weeks	01-06-2025
Vitamin D Supplement	1 X Per Week 4 Weeks	01-06-2025

Patient consultation notes will appear here...

Click the microphone to start recording.

Hello! I can help you with patient notes, medical queries, create personalized care pathways and other documentation. How can I assist you today?

Fig 8.1 Clinician dashboard main screen

8.2 Visualizing behavioral Data

Alongside clinical information, clinicians can review detailed behavioral information by clicking on the “Behavioral Data” tab. This section presents a visual summary of patient behavior type, their quality-of-life-based goals, symptom trends, and mood trends. This information can be captured from patient facing application as described in the previous chapter 7.4.

Previous research from co-create sessions and semi-structured interviews show that clinicians need structured ways to act on behavioral signals. The dashboard presents patient goals, adherence patterns, activation levels, and mood trends all in one place.

Each patient’s PAM score appears alongside behavioral trends to help clinicians track engagement and tailor follow-up communication and other activities. Color-coded badges indicate behavior types such as proactive, reactive, or passive giving the clinicians a quick view of what to expect from patients.

The dashboard supports asynchronous review of a patient’s goals. Clinicians can review behavioral insights ahead of appointments or during care plan adjustments. This improves continuity without increasing the workload. It also enables earlier intervention when patients begin to disengage.

Mood trends and symptom logs are summarized visually to flag emotional risk. This helps clinicians adapt treatment tone or initiate referrals to psychological support.

By placing behavior and emotion alongside clinical information, the interface bridges the gap between patient experience and medical action. It also enables clinicians to filter patients by activation level and behavior type, helping them prioritize outreach. This leads to better adherence, stronger trust, and more focused use of consultation time.

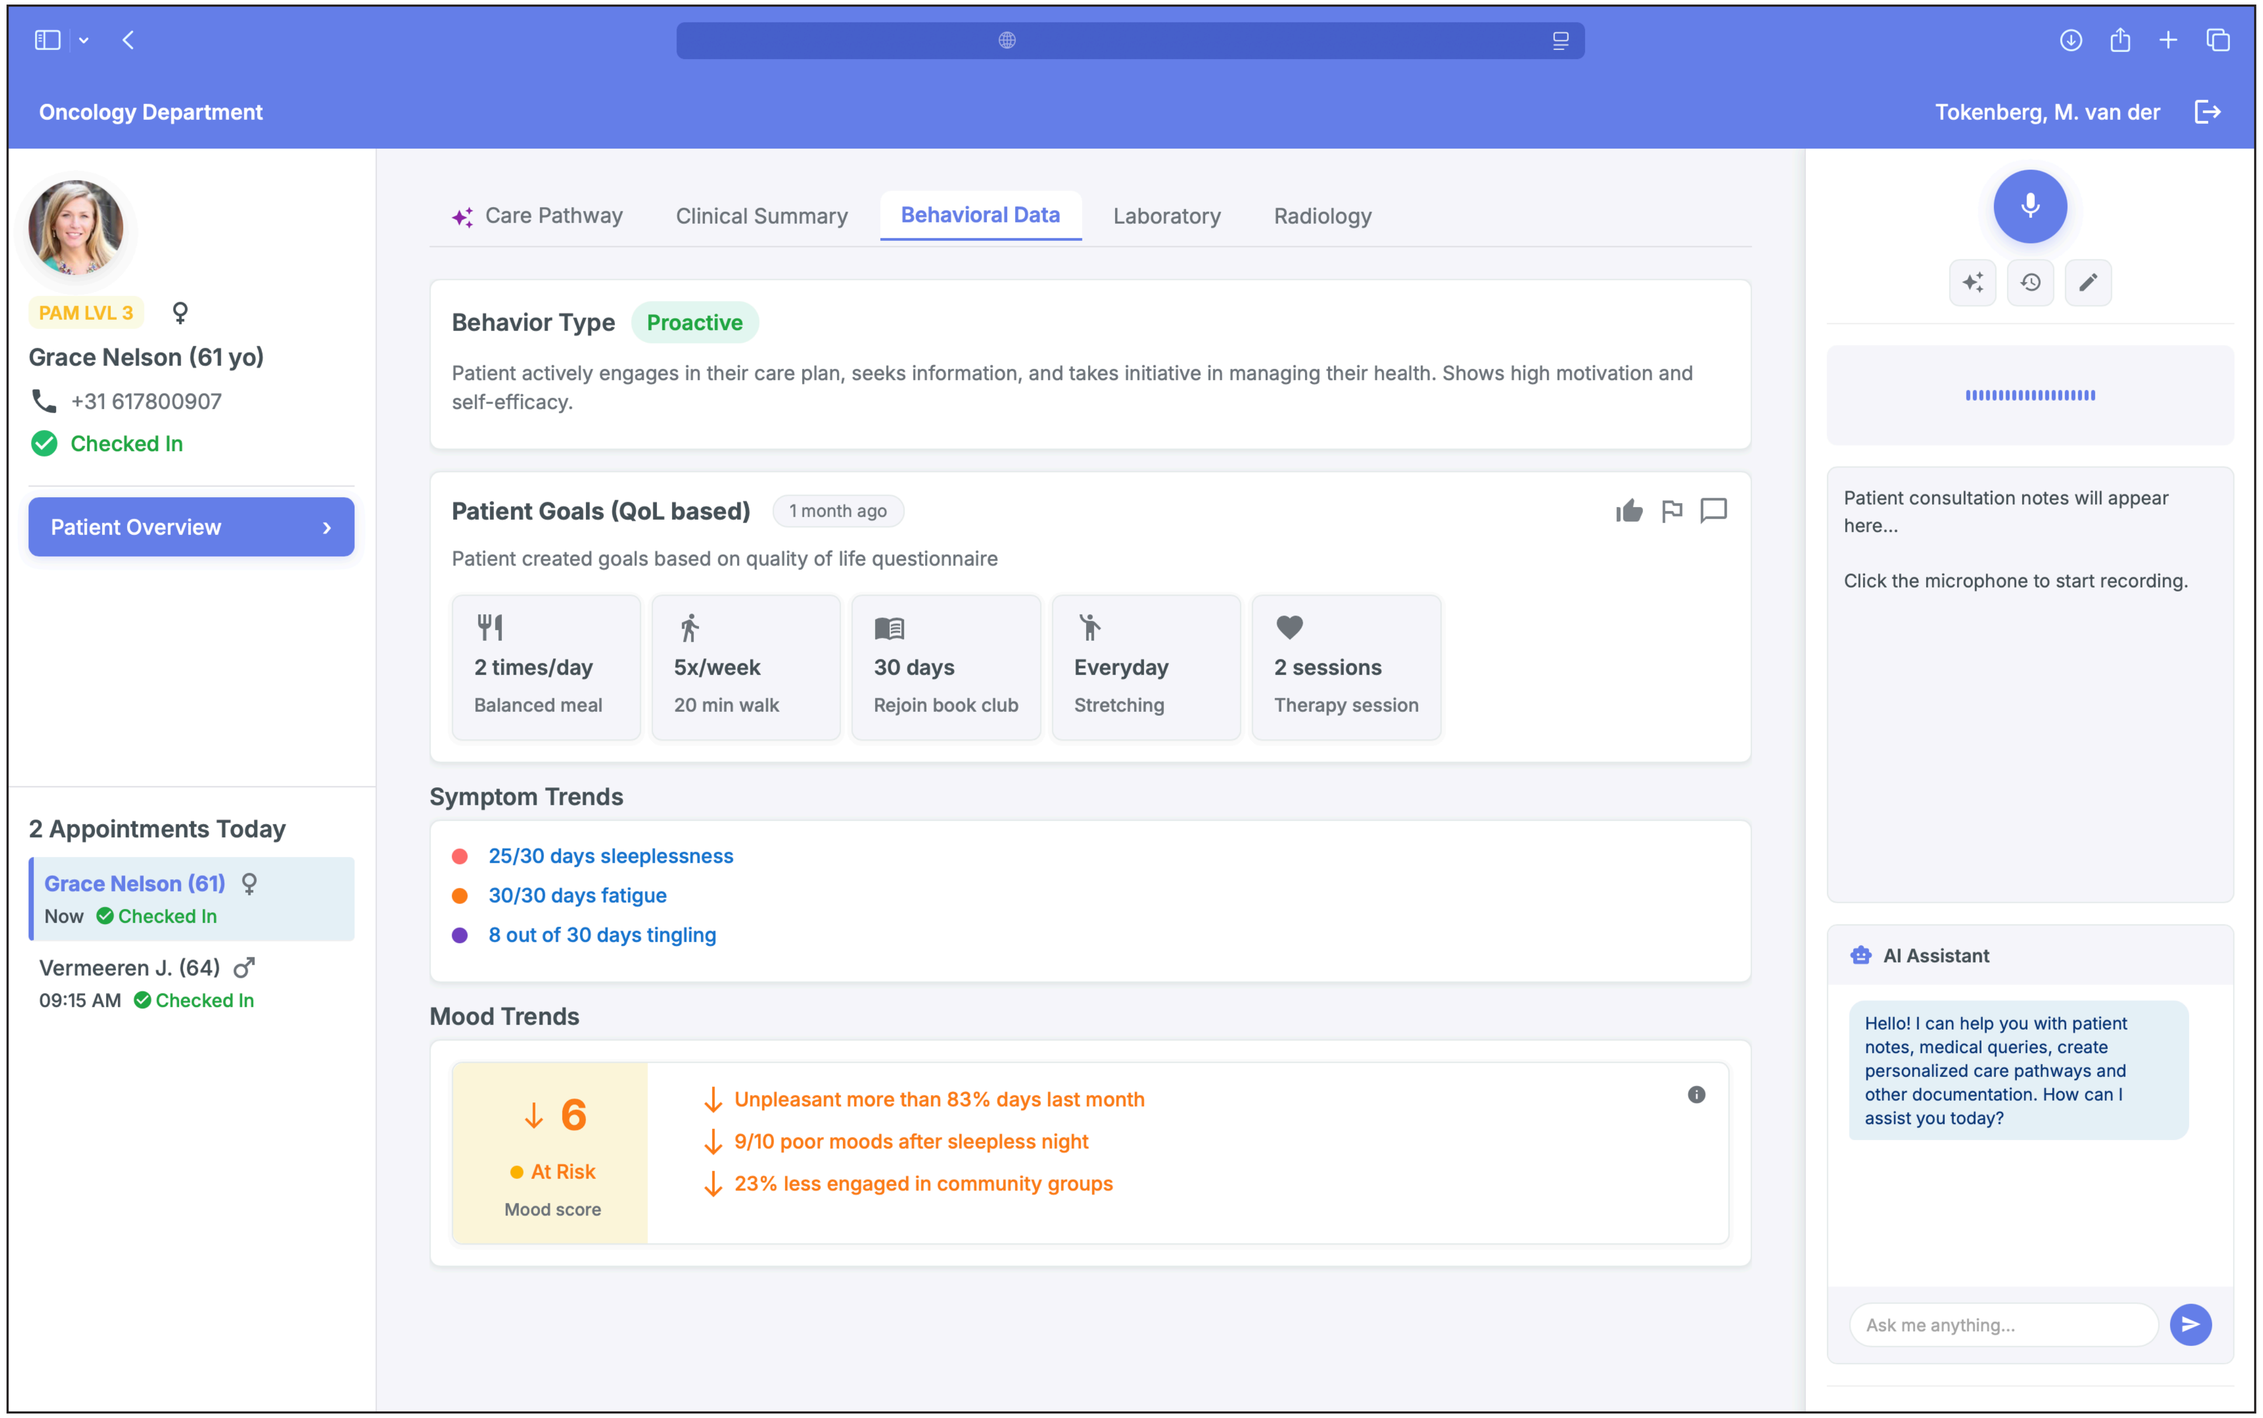

Grace Nelson (61 yo)
+31 617800907
Checked In

2 Appointments Today
Grace Nelson (61) ♀
Now ✓ Checked In
Vermeeren J. (64) ♂
09:15 AM ✓ Checked In

Behavior Type Proactive

Patient actively engages in their care plan, seeks information, and takes initiative in managing their health. Shows high motivation and self-efficacy.

Patient Goals (QoL based) 1 month ago

Patient created goals based on quality of life questionnaire

2 times/day	5x/week	30 days	Everyday	2 sessions
Balanced meal	20 min walk	Rejoin book club	Stretching	Therapy session

Symptom Trends

- 25/30 days sleeplessness
- 30/30 days fatigue
- 8 out of 30 days tingling

Mood Trends

↓ 6
At Risk
Mood score

↓ Unpleasant more than 83% days last month
↓ 9/10 poor moods after sleepless night
↓ 23% less engaged in community groups

Patient consultation notes will appear here...

Click the microphone to start recording.

AI Assistant

Hello! I can help you with patient notes, medical queries, create personalized care pathways and other documentation. How can I assist you today?

Ask me anything...

Fig 8.2 Behavioral data visualized on the dashboard

8.3 Agent-led Personalization of Care Pathways

Recalling the concerns from earlier chapters specifically chapter 1 section 1.2, one of the significant challenges was the creation of these pathways being resource-intensive, requiring significant time, effort, and dedicated personnel. It also restricted doctor's autonomy in making decisions since doctors often felt like their hands were tied while putting care pathways into practice (Khalifa & Alswailem, 2015).

In this section, I introduce a new concept: agent-led personalization of care pathways. In this context, an agent is an intelligent assistant that can observe patient data, interpret it, and suggest care actions to the clinician. A helpful way to think about this is to compare the agent to Google Maps. You enter your destination, and the app processes live traffic data, road closures, and your current location to suggest the fastest or safest route. You remain in control of the car, but the app helps you avoid delays and make better decisions faster. Similarly, the agent reviews patient data, highlights relevant care actions, and leaves the final decisions to the clinician.

By treating the interface as an intelligent agent rather than just a tool, the dashboard supports clinicians in personalizing care without increasing administrative burden. The intelligent system processes multiple input sources (clinical data, behavioral data, consultation notes, etc) and offers recommendations that clinicians can trust and adapt as needed. This interaction reflects a collaborative model where the AI proposes options and clinicians remain in control of final decisions. This significantly reduces the time clinicians spend piecing together information and translating it into action.

The interface presents the personalized care plan as a set of modular recommendations (as shown in fig 8.3). Each care module includes a title, a short description, relevant clinical and behavioral inputs, and reference links to established guidelines. These modules bring evidence-based practices into the clinician's workflow without requiring manual lookup.
Clinicians can approve, reject, or edit each module. The interface allows for in-line editing and immediate visual feedback on module status. This supports flexibility and maintains clinical autonomy. To support team-based care, the system automatically suggests resource assignments for each module (see fig 8.4).

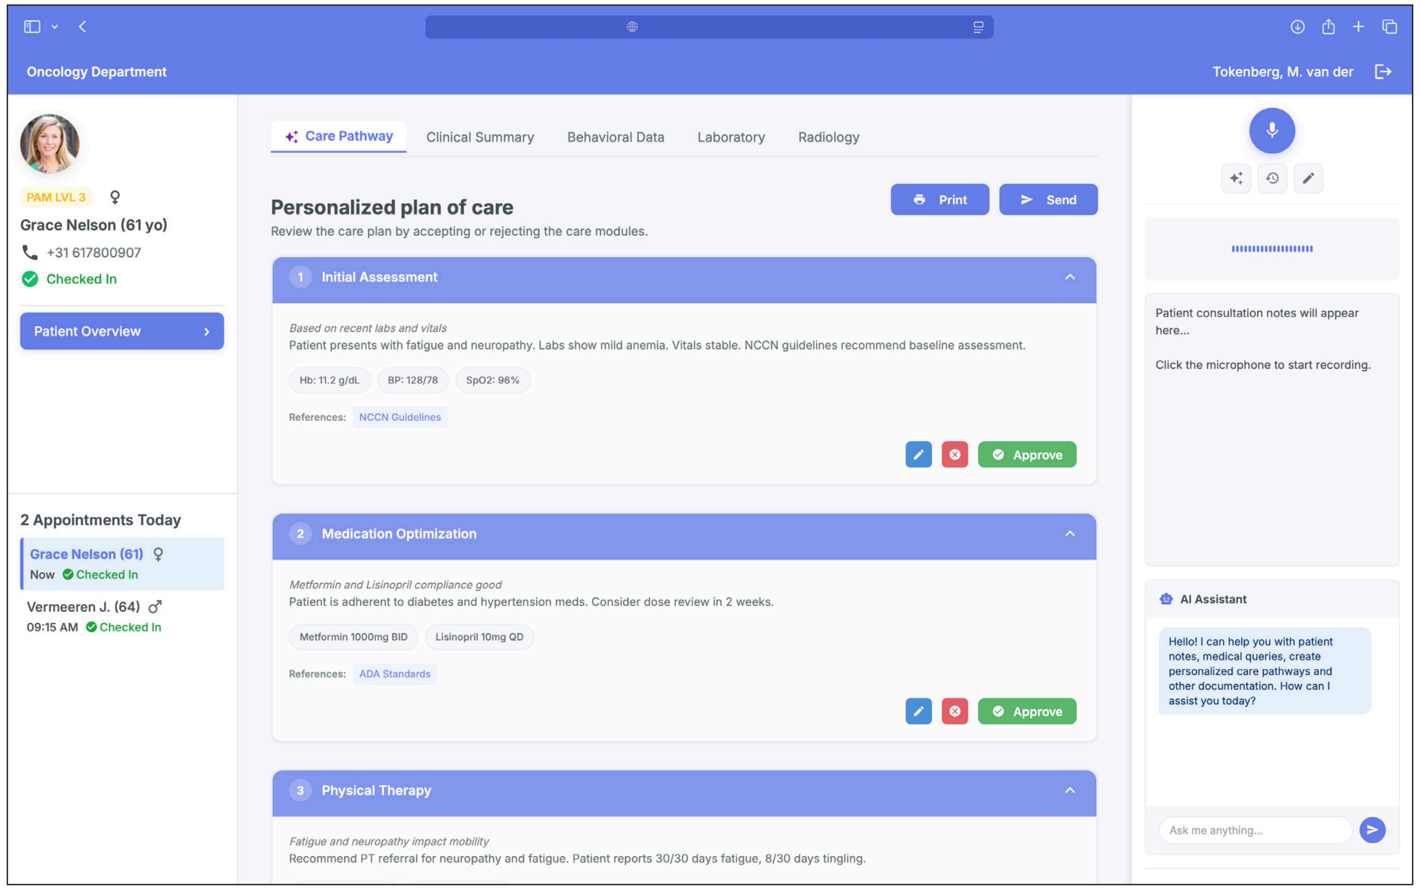

Oncology Department
Tokenberg, M. van der

PAM LVL 3
Grace Nelson (61 yo)
+31 617800907
Checked In
Patient Overview

2 Appointments Today
Grace Nelson (61)
Now Checked In
Vermeeren J. (64)
09:15 AM Checked In

Personalized plan of care
Review the care plan by accepting or rejecting the care modules.

1 Initial Assessment
Based on recent labs and vitals
Patient presents with fatigue and neuropathy. Labs show mild anemia. Vitals stable. NCCN guidelines recommend baseline assessment.
Hb: 11.2 g/dL
BP: 128/78
SpO2: 96%
References: NCCN Guidelines

2 Medication Optimization
Metformin and Lisinopril compliance good
Patient is adherent to diabetes and hypertension meds. Consider dose review in 2 weeks.
Metformin 1000mg BID
Lisinopril 10mg QD
References: ADA Standards

3 Physical Therapy
Fatigue and neuropathy impact mobility
Recommend PT referral for neuropathy and fatigue. Patient reports 30/30 days fatigue, 8/30 days tingling.

Patient consultation notes will appear here...

Click the microphone to start recording.

AI Assistant
Hello! I can help you with patient notes, medical queries, create personalized care pathways and other documentation. How can I assist you today?
Ask me anything...

Fig 8.3 Personalized care plan as a set of care modules

It displays the assigned clinician's name for specialized care requirements. This ensures that each recommendation flows to the right specialist, from nutritionists and psychologists to oncology nurses and general practitioners.

Once finalized, the interface automatically generates a patient-facing version of the care plan in plain language. This feature stems out of a user pain point discussed in chapter 7 section 7.1 in which patients find it difficult to navigate technical jargon and often feel anxious. This version of the care plan is designed to specifically avoid clinical jargon and break the recommendations into simple, step-by-step instructions. It uses encouraging language and contact details for each assigned provider. It also includes reasoning behind the administered medications creating transparency in care delivery. This improves trust, strengthens patient engagement, and helps close the loop between decision and communication.

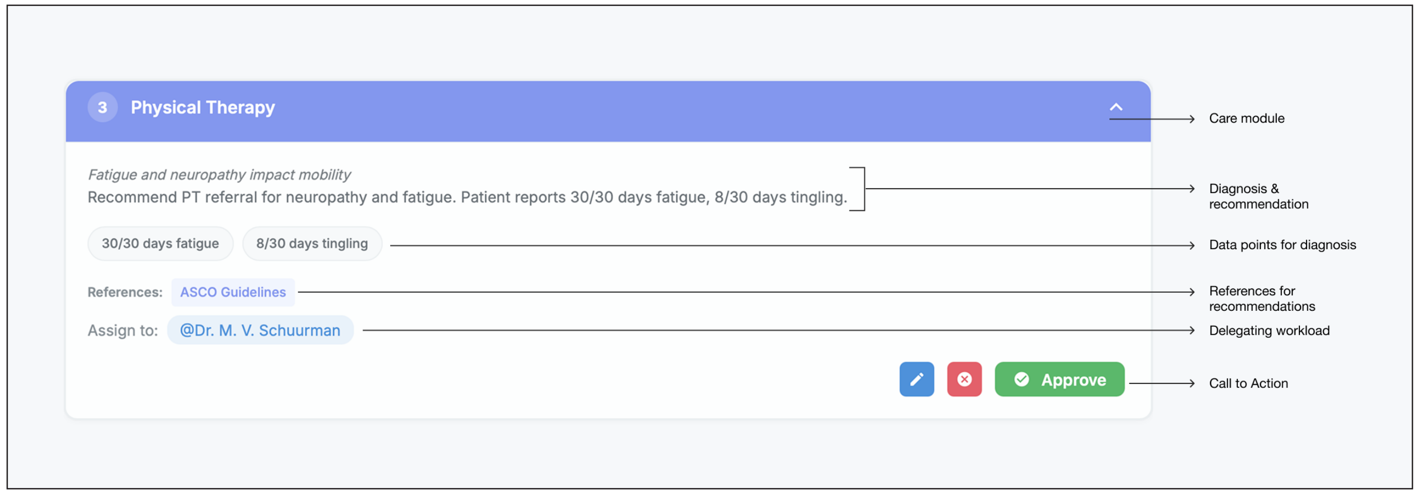

Fatigue and neuropathy impact mobility
Recommend PT referral for neuropathy and fatigue. Patient reports 30/30 days fatigue, 8/30 days tingling.
References: ASCO Guidelines
Assign to: @Dr. M. V. Schuurman

Fig 8.4 Care module structure

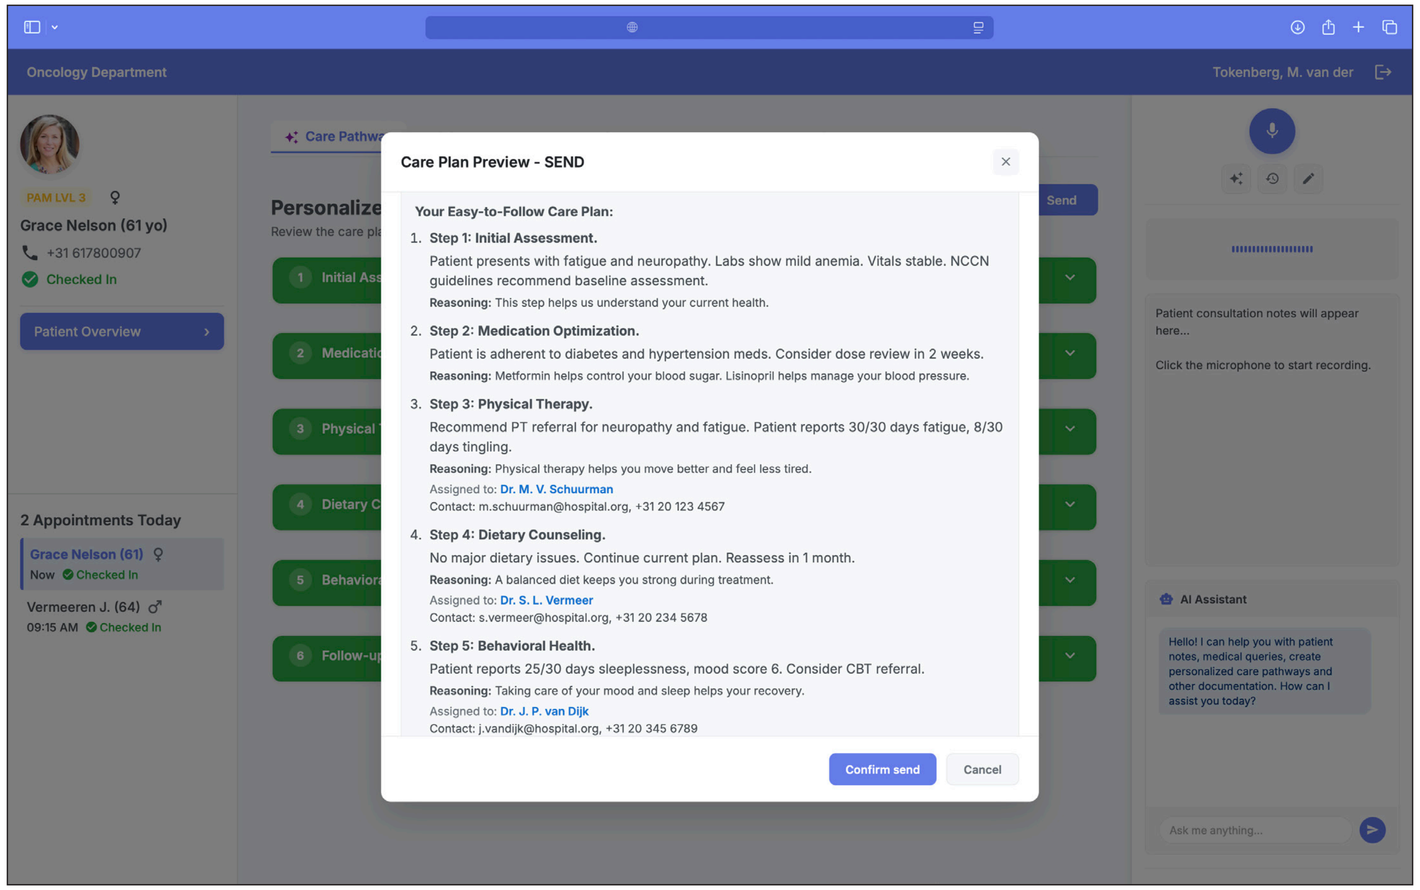

Care Plan Preview - SEND

Your Easy-to-Follow Care Plan:

1. **Step 1: Initial Assessment.**
   Patient presents with fatigue and neuropathy. Labs show mild anemia. Vitals stable. NCCN guidelines recommend baseline assessment.
   Reasoning: This step helps us understand your current health.
2. **Step 2: Medication Optimization.**
   Patient is adherent to diabetes and hypertension meds. Consider dose review in 2 weeks.
   Reasoning: Metformin helps control your blood sugar. Lisinopril helps manage your blood pressure.
3. **Step 3: Physical Therapy.**
   Recommend PT referral for neuropathy and fatigue. Patient reports 30/30 days fatigue, 8/30 days tingling.
   Reasoning: Physical therapy helps you move better and feel less tired.
   Assigned to: **Dr. M. V. Schuurman**
   Contact: m.schuurman@hospital.org, +31 20 123 4567
4. **Step 4: Dietary Counseling.**
   No major dietary issues. Continue current plan. Reassess in 1 month.
   Reasoning: A balanced diet keeps you strong during treatment.
   Assigned to: **Dr. S. L. Vermeer**
   Contact: s.vermeer@hospital.org, +31 20 234 5678
5. **Step 5: Behavioral Health.**
   Patient reports 25/30 days sleeplessness, mood score 6. Consider CBT referral.
   Reasoning: Taking care of your mood and sleep helps your recovery.
   Assigned to: **Dr. J. P. van Dijk**
   Contact: j.vandijk@hospital.org, +31 20 345 6789

Fig 8.5 Patient facing care plan avoiding technical jargon

8.4 Prototyping workflow using Cursor AI

The interfaces shown above was built using a rapid design and development approach. I chose Cursor AI as the primary development environment because it allows rapid prototyping with AI-assisted coding. I also had prior experience using Cursor to build and deploy interfaces to platforms like GitHub and Vercel, which gave me confidence in maintaining and modifying the codebase.

My main goal was to build a high-fidelity prototype quickly, so I could test it with real clinicians. From experience, I knew that most doctors respond better to functional tools than abstract ideas. They need to click through a design and imagine it in their workflow before giving actionable feedback. Cursor allowed me to translate my Figma designs into working HTML and CSS components at speed.

To ensure visual consistency with hospital systems, I referred to screenshots of the HIX platform and used them as a base layout for the prototype. This decision helped preserve the mental models of clinical users and created a sense of familiarity. I initially focused on surfacing behavioral data next to clinical information, but during a review session, my supervisor pointed out that I was still placing the burden of personalization on the clinician. This pushed me to extend the concept into an agentic interface that could actively assist with care planning.

The most challenging part of the design was defining the content structure of the care modules. I spent a lot of time selecting framing the information in a way that supported evidence-based practice. I also went through several iterations of button placement and labelling to reduce friction. For example, the approve and assign features went through multiple rounds of informal testing with peers from TU Delft Industrial Design Engineering (IDE).

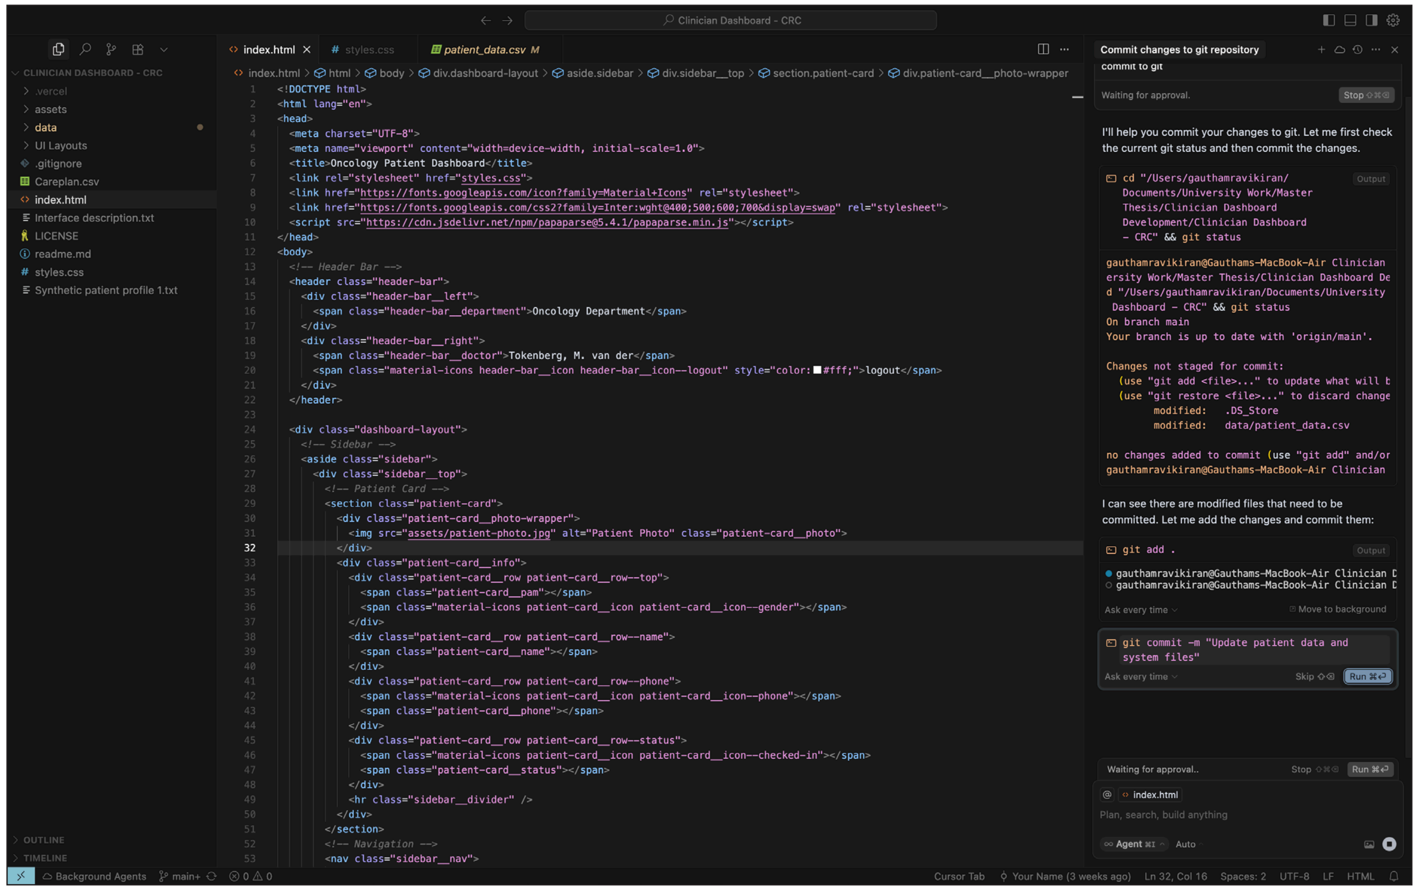

```html
<!DOCTYPE html>
<html lang="en">
<head>
  <meta charset="UTF-8">
  <meta name="viewport" content="width=device-width, initial-scale=1.0">
  <title>Oncology Patient Dashboard</title>
  <link rel="stylesheet" href="styles.css">
  <link href="https://fonts.googleapis.com/icon?family=Material+Icons" rel="stylesheet">
  <link href="https://fonts.googleapis.com/css2?family=Inter:wght@400;500;600;700&display=swap" rel="stylesheet">
  <script src="https://cdn.jsdelivr.net/npm/papaparse@5.4.1/papaparse.min.js"></script>
</head>
<body>
  <!-- Header Bar -->
  <header class="header-bar">
    <div class="header-bar__left">
      <span class="header-bar__department">Oncology Department</span>
    </div>
    <div class="header-bar__right">
      <span class="header-bar__doctor">Tokenberg, M. van der</span>
      <span class="material-icons header-bar__icon header-bar__icon--logout" style="color:#fff;">logout</span>
    </div>
  </header>

  <div class="dashboard-layout">
    <!-- Sidebar -->
    <aside class="sidebar">
      <div class="sidebar__top">
        <!-- Patient Card -->
        <section class="patient-card">
          <div class="patient-card__photo-wrapper">
            <img src="assets/patient-photo.jpg" alt="Patient Photo" class="patient-card__photo">
          </div>
          <div class="patient-card__info">
            <div class="patient-card__row patient-card__row--top">
              <span class="patient-card__pam"></span>
              <span class="material-icons patient-card__icon patient-card__icon--gender"></span>
            </div>
            <div class="patient-card__row patient-card__row--name">
              <span class="patient-card__name"></span>
            </div>
            <div class="patient-card__row patient-card__row--phone">
              <span class="material-icons patient-card__icon patient-card__icon--phone"></span>
              <span class="patient-card__phone"></span>
            </div>
            <div class="patient-card__row patient-card__row--status">
              <span class="material-icons patient-card__icon patient-card__icon--checked-in"></span>
              <span class="patient-card__status"></span>
            </div>
            <hr class="sidebar__divider" />
          </div>
        </section>
        <!-- Navigation -->
        <nav class="sidebar__nav">
```

Commit changes to git repository

commit to git

Waiting for approval.

I'll help you commit your changes to git. Let me first check the current git status and then commit the changes.

```
cd "/Users/gauthamravikiran/
Documents/University Work/Master
Thesis/Clinician Dashboard
Development/Clinician Dashboard
- CRC" && git status

gauthamravikiran@Gauthams-MacBook-Air Clinician
ersity Work/Master Thesis/Clinician Dashboard De
d "/Users/gauthamravikiran/Documents/University
 Dashboard - CRC" && git status
On branch main
Your branch is up to date with 'origin/main'.

Changes not staged for commit:
  (use "git add <file>..." to update what will b
  (use "git restore <file>..." to discard change
        modified:   .DS_Store
        modified:   data/patient_data.csv

no changes added to commit (use "git add" and/or
gauthamravikiran@Gauthams-MacBook-Air Clinician
```

I can see there are modified files that need to be committed. Let me add the changes and commit them:

```
git add .

● gauthamravikiran@Gauthams-MacBook-Air Clinician D
○ gauthamravikiran@Gauthams-MacBook-Air Clinician D
```

```
git commit -m "Update patient data and
system files"
```

Fig 8.6 Cursor AI development environment

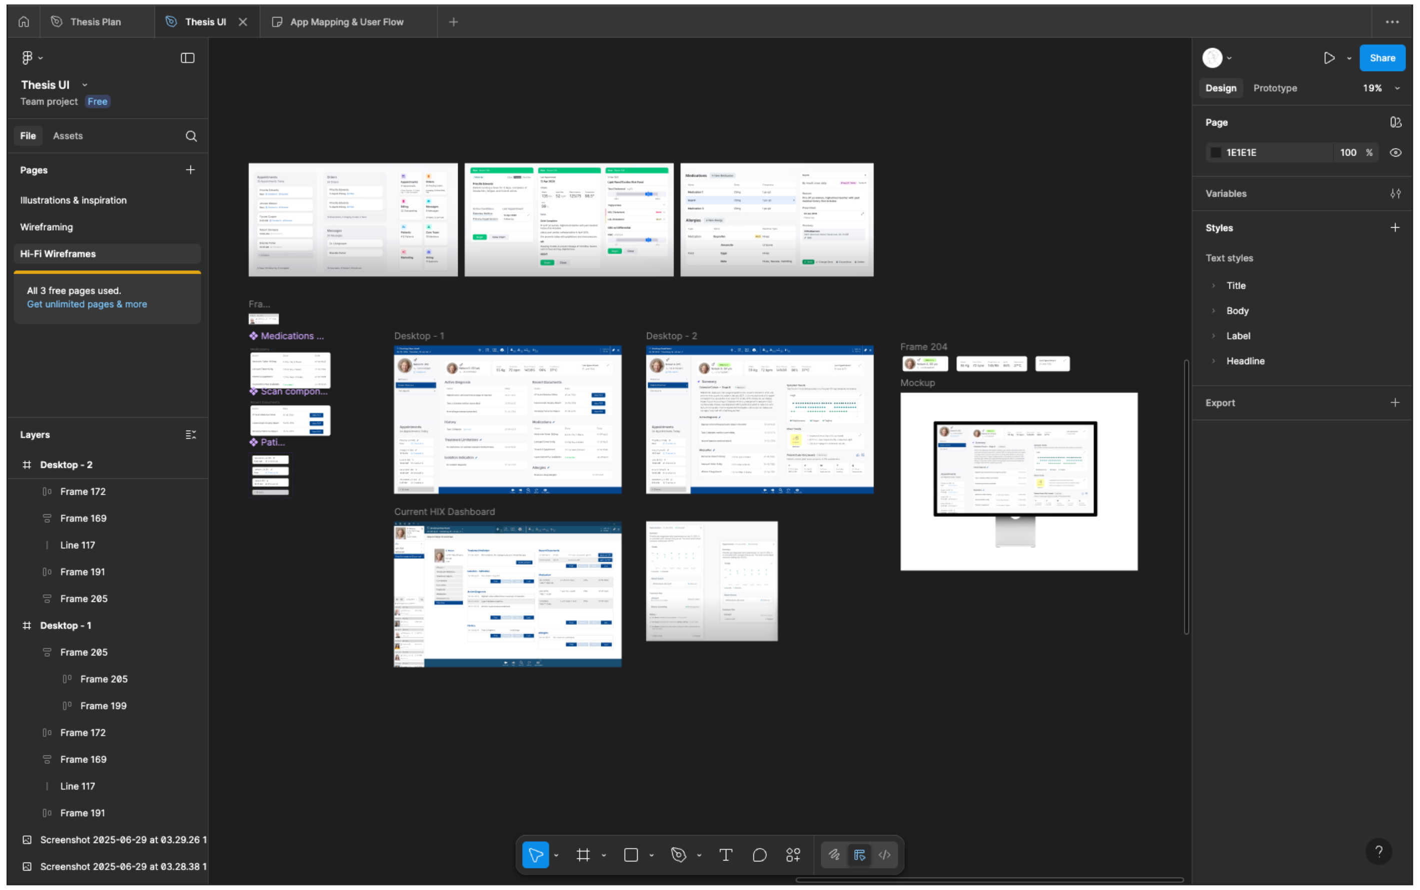

Fig 8.7 Figma workspace for initial UI design

8.5 Feedback in the wild: A chance encounter that changed the interface

Before developing this high-fidelity concept, I had an opportunity to attend the ASPIRE Conference on AI in healthcare in Hong Kong. During my visit, I met a master's student in Clinical Technology who was also from TU Delft. She was fully familiar with the HIX platform used across hospitals in the Netherlands.

After I shared some details about my thesis, she expressed interest in trying the dashboard prototype. Though the interface was not fully complete, I allowed her to explore it. The feedback she provided was rich with insights that significantly shaped the final design.

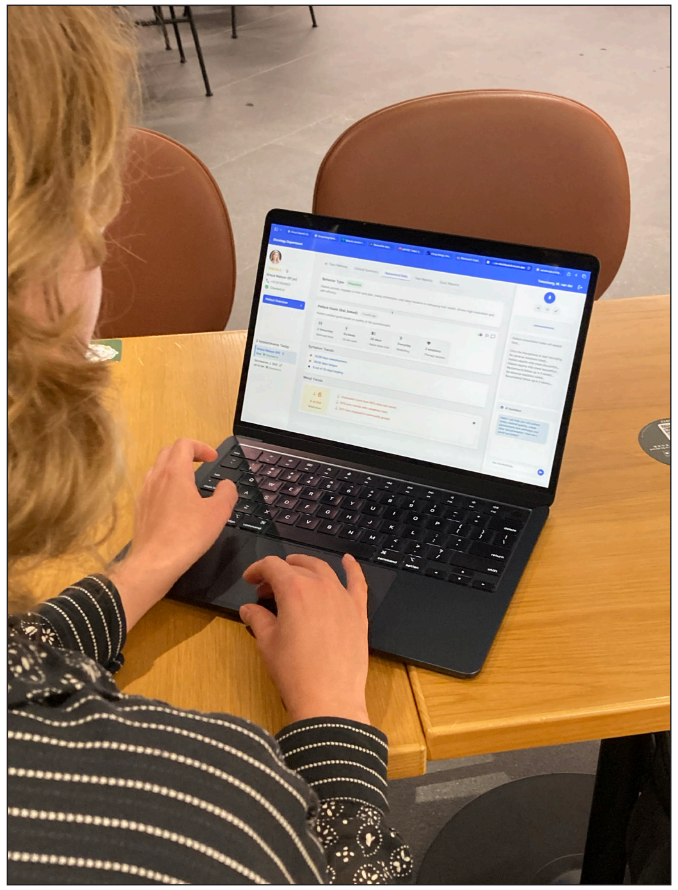

Fig 8.8 Informal feedback from a Clinical Technology Master student from TU Delft

She responded positively to the care module reference feature and found it helpful for clinical credibility. She also emphasized that doctors often hold multiple roles and said that assigning care modules to the right person within the interface would save time and reduce confusion. Based on this input, I implemented the doctor assignment feature described in earlier sections, that matches modules to roles such as nurses, psychologists, or general practitioners.

She also suggested changing several tab names to align better with clinical terminology. For example, she recommended renaming “Scan results” to “Radiology” and “Test results” to “Laboratory,” as these terms are more familiar to clinicians. I integrated these changes into the interface.

During our discussion, she also highlighted how valuable a quick patient summary would be. She explained that current systems like HIX include a feature called “Naaslag,”(dutch word for reference point). This contains semi-structured notes from multiple clinicians who have seen the patient in the past. Clinicians must scroll through long histories to understand the patient's case. The concise patient summary in my design helped reduce this friction and was well received.

She strongly advocated for a plain-language care plan that patients could understand. Drawing on her feedback, I built in a feature that translates clinical care plans into accessible language. It breaks down the treatment into steps, avoids jargon, and includes explanations for medications. This aligns with findings from Chapter 7 and adds transparency to care delivery.

This informal test confirmed many of my design decisions and introduced new features that made the prototype more grounded and usable. It showed how spontaneous feedback can offer rich insight and bring clarity to user-centered design.

8.6 Reflections

This chapter taught me that design is not just about tools and visuals, but about timing, openness, and readiness. Who would have thought that I would meet someone during an international conference who would take such a deep interest in my work and help shape it so meaningfully? That conversation reminded me to always be prepared to share my work, even if it feels imperfect. The unexpected user test became one of the most valuable feedback moments in this thesis. It brought in a perspective from someone deeply familiar with the clinical systems I was designing for. Her insights helped transform my interface from a static tool into a dynamic, usable, and clinically relevant prototype. It validated my direction while also challenging me to improve weak spots.

From a methodological perspective, this phase taught me how to use design as a tool for inquiry. I prototyped to think, to ask, and to provoke feedback. Cursor helped me move quickly, but the true shift came from being willing to let go of assumptions & listen to users.

The idea of agent-led personalization became clearer through these iterations. I began to see AI not as a replacement for clinicians but as a partner in decision making. Like a map that shows the best route but leaves you in control, the interface supported clinicians without taking over. It highlighted important actions, respected their judgment, and focused on saving time while keeping care thoughtful.

This experience showed me that careful design and timely feedback can transform both the design and the way we think about delivering care.

8.7 Conclusion

This chapter began with three guiding research questions. How can behavioral data be visualized to support clinical decision making? How might we reduce the burden of personalizing care from a clinician's perspective? How can AI support personalization while preserving clinician autonomy?

Through the design and prototyping of a clinician-facing dashboard, I explored each of these questions in practice. The result was a future-facing interface that integrates behavioral insights, clinical data, and intelligent decision support into a single experience. Rather than designing yet another data viewer, this chapter focused on creating an agentic interface that could actively assist the clinician in making sense of patient behavior and translating that into actionable care plans.

One of the key strengths of the dashboard lies in how it makes behavioral data visible and useful. The system surfaces PAM scores, patient goals, and emotional trends in clear and accessible ways. It allows clinicians to scan for disengagement, flag risk, and prepare for more empathic and effective conversations. These insights are no longer hidden in long text notes or separate systems. They are placed alongside clinical metrics where decisions happen.

The agent-led care pathway generator pushed this idea further. It showed how we can use AI not just to predict, but to support. The agent gathers signals from clinical records, behavioral data, and consultation history. It then suggests care modules that reflect both evidence-based medicine and personalized need. These modules are not final answers but are prompts for action. Clinicians retain the authority to edit, approve, or reject the suggestion. This approach preserves autonomy while reducing cognitive and administrative load.

Perhaps the most transformative part of this chapter was the unplanned user test that took place during the ASPIRE conference. The informal conversation with a clinical technology student brought a fresh wave of insight. It validated design choices and led to specific improvements, including the doctor assignment feature and clearer clinical language across tabs. This moment reminded me that good design happens not only in structured phases, but also through spontaneous dialogue. Being ready to share unfinished work opened the door to unexpected clarity.

The goal of this exercise was to build a functional prototype that makes the future of care planning easier to explore, question, and improve. The dashboard served as a design probe that helped me reflect on what it means to support clinicians without overwhelming them. It also gave form to the idea that AI can collaborate rather than control.

In the next chapter, I test this interface with expert users through guided walkthroughs and interviews. Their feedback provides deeper insight into the practicality, usability, and acceptance of agent-led care planning in real clinical environments.

Chapter 9
Expert Feedback session

This chapter presents the expert validation phase of the project. Its purpose is to evaluate how clinicians interpret and respond to the Vision 2040 roadmap and the clinician dashboard prototype introduced in Chapter 8. This stage focuses on validating the strategic direction of the project and assessing the clinical relevance of the interface in real-world settings.

The aim was to investigate how structured behavioral data and AI-assisted tools could be integrated into routine workflows without increasing clinician burden. This step grounds the design vision in clinical reality and tests whether the system aligns with the values, capacities, and decision-making processes of its intended users. The following research questions guided this inquiry:

“What is the impact of providing structured behavioral data alongside clinical patient information?”

“How do clinicians view AI assisted, behavior aware decision support tools in practice?”

To explore these questions, I conducted a series of expert walkthrough sessions with medical professionals at Erasmus MC. Each session included a short presentation on the patient profiling process and the Vision 2040 roadmap. After this, participants interacted with a high-fidelity prototype of the clinician dashboard. During the session, I observed how participants navigated the interface and collected open-ended feedback on specific features and broader concepts.

This chapter presents findings drawn exclusively from these structured and recorded feedback sessions. The expert insights shared here provide critical validation of the system’s direction and answers the central research question of how data-enabled design methods may be used to personalize care pathways while reducing the cognitive and administrative load on healthcare professionals.

9.1 Conducting the sessions

To evaluate the clinical relevance and future potential of the Vision 2040 roadmap and the AI-powered clinician dashboard, I conducted three expert interviews with healthcare professionals at Erasmus MC. These interviews were held in person and organized through personal networks developed during the co-creation phase described in Chapter 5. All three participants had previously taken part in the co-design sessions and were already familiar with the project’s focus and objectives.

Fig 9.1 Erasmus Medical Center

The participants included one PhD researcher in surgical oncology, one PhD researcher in Pulmonary medicine, and the Head of Oncological and Gastrointestinal Surgery at Erasmus MC. Each brought a unique clinical perspective informed by direct experience with chronic care coordination and long-term patient follow-up.

The sessions followed a semi-structured interview format. Each began with a brief presentation summarizing earlier findings from the patient-facing application and profiling process. This was followed by a walkthrough of the Vision 2040 roadmap, used to contextualize the long-term ambition of the project. The central activity was a guided interaction with the high-fidelity prototype of the clinician dashboard. Participants navigated key features such as the behavioral insights view, care module suggestions, and patient-facing care plan output.

Although parts of the discussion touched on patient data collection methods, the primary objective was to examine the dashboard interface and strategic direction from a healthcare professional’s perspective. The focus remained on how these tools could integrate into clinical workflows, support decision-making, and impact future models of care delivery. In the next section, we discuss the key themes that emerged during these sessions, with a focus on clinician attitudes, concerns, and recommendations related to the proposed system.

9.2 Findings from the sessions

The expert interviews revealed recurring themes across all participants, centered on the clinical value, usability, and strategic feasibility of behavior-aware and AI-supported decision-making tools. This section presents the main findings under six thematic categories. Each theme integrates expert observations, concerns, and suggestions related to the dashboard prototype and the Vision 2040 roadmap.

Behavioral Data enhances clinical understanding

All experts acknowledged the value of integrating structured behavioral data into clinical workflows. They emphasized that such data adds a dimension of patient understanding that is often known intuitively but rarely formalized. These insights were viewed as particularly useful during pre-consultation review.

Experts appreciated that the dashboard surfaced behavioral traits such as proactiveness, emotional distress, and goal orientation through visualizations of PAMs. These attributes were often discussed in multidisciplinary meetings but were not documented in a structured format. But at the same time, clinicians reiterated that behavioral traits evolve over time. They stressed the importance of combining algorithmic behavioral summaries with long-term clinical intuition.

> “I would definitely look at behavioral data before consultation. That might help shape the conversation. It would also be valuable to have some input from the clinician side, because they get to know a patient over the course of their journey. That’s something you experience through conversation.”
>
> *- PhD researcher in surgical oncology*

AI suggestions must support, not replace, decision making

The AI-generated care modules received strong positive feedback. Experts highlighted that the system’s design struck the right balance by providing guidance without prescribing fixed solutions. The transparency of AI reasoning and ability to trace recommendations back to known clinical sources was seen as essential.

> “It gives suggestions, not decisions. That’s important. Just giving information or suggestions makes it a non-MDR regulated, non-medical device. So you stay outside that scope.”
>
> *- PhD researcher in surgical oncology*

Experts saw the care module generator as a new way to adapt treatment plans considering behavioral patterns and patient lifestyle.

> “This is really something new, a basis to know what patients do in daily life and adjust treatment or motivate them to walk more. It really helps clinicians”
>
> *- PhD researcher in Pulmonary medicine*

They also supported the connection between AI logic and guideline-based care, recommending the integration of national clinical sources. But clinicians also stressed that autonomy must remain with the physician, and the system should never position its outputs as definitive answers.

The most enthusiastic endorsement came from the Head of Oncology, who emphasized the practical value of embedding such a tool directly into existing systems

Fig 9.2 Clinician exploring the clinical and behavioral summary of the dashboard

> “It’s tremendous work actually. If this could be embedded in our electronic patient file, that would really be a big help.”
>
> *- Head of Oncological and Gastrointestinal Surgery on AI based personalization feature*

Operational burden must be minimized

A recurring concern was the risk of adding to the workload of already time-constrained clinicians. Experts were cautious about the asynchronous input and goal-review features, especially if they introduced unpredictability or required constant monitoring.

> “I’m a little bit anxious about the amount of work it might take for a doctor. You would need to build in specific moments to check up on these things or make it an administrative consultation
>
> We want to prevent that these things would cause a lot more work. You’d have to check every patient who filled something in not knowing when they did.”

- PhD researcher in surgical oncology

Experts recommended tying goal review into structured consultation moments. They emphasized that new tools must fit within current workflows and time blocks to avoid reactive interruptions.

Shared access empowers patients and improves health literacy

All experts endorsed the notion of behavioral data being shared with patients as a reflection tool. They saw potential in dual-purpose forms that inform both patients and clinicians.

Participants strongly supported the use of layered educational models that offer basic explanations first, with options to go into deeper, verified content such as kanker.nl.

> “These should not only be filled in for the clinician, but also for the patients themselves, so they can track sleep, mood, whatever. Patients would be able to see their own scores and track them over time… maybe give some online seminars on sleep.
>
> We’re working a lot on layering information. First layer is the basic version; then you can go deeper if needed. And finally, we refer to kanker.nl”

- PhD researcher in surgical oncology

Structured overviews are essential

The clinician dashboard was consistently praised for its clear structure and calm visual layout. Experts compared it favorably to current systems like HIX, which they described as overwhelming and repetitive.

> “I get tired looking at HIX all day. This is calmer. To keep the overview, this is very good - to have a summary of what the care pathway looks like”

- PhD researcher in Pulmonary medicine

Participants criticized the need to scroll through long histories or repetitive entries in existing systems, noting that the dashboard’s summarization model provided a meaningful improvement.

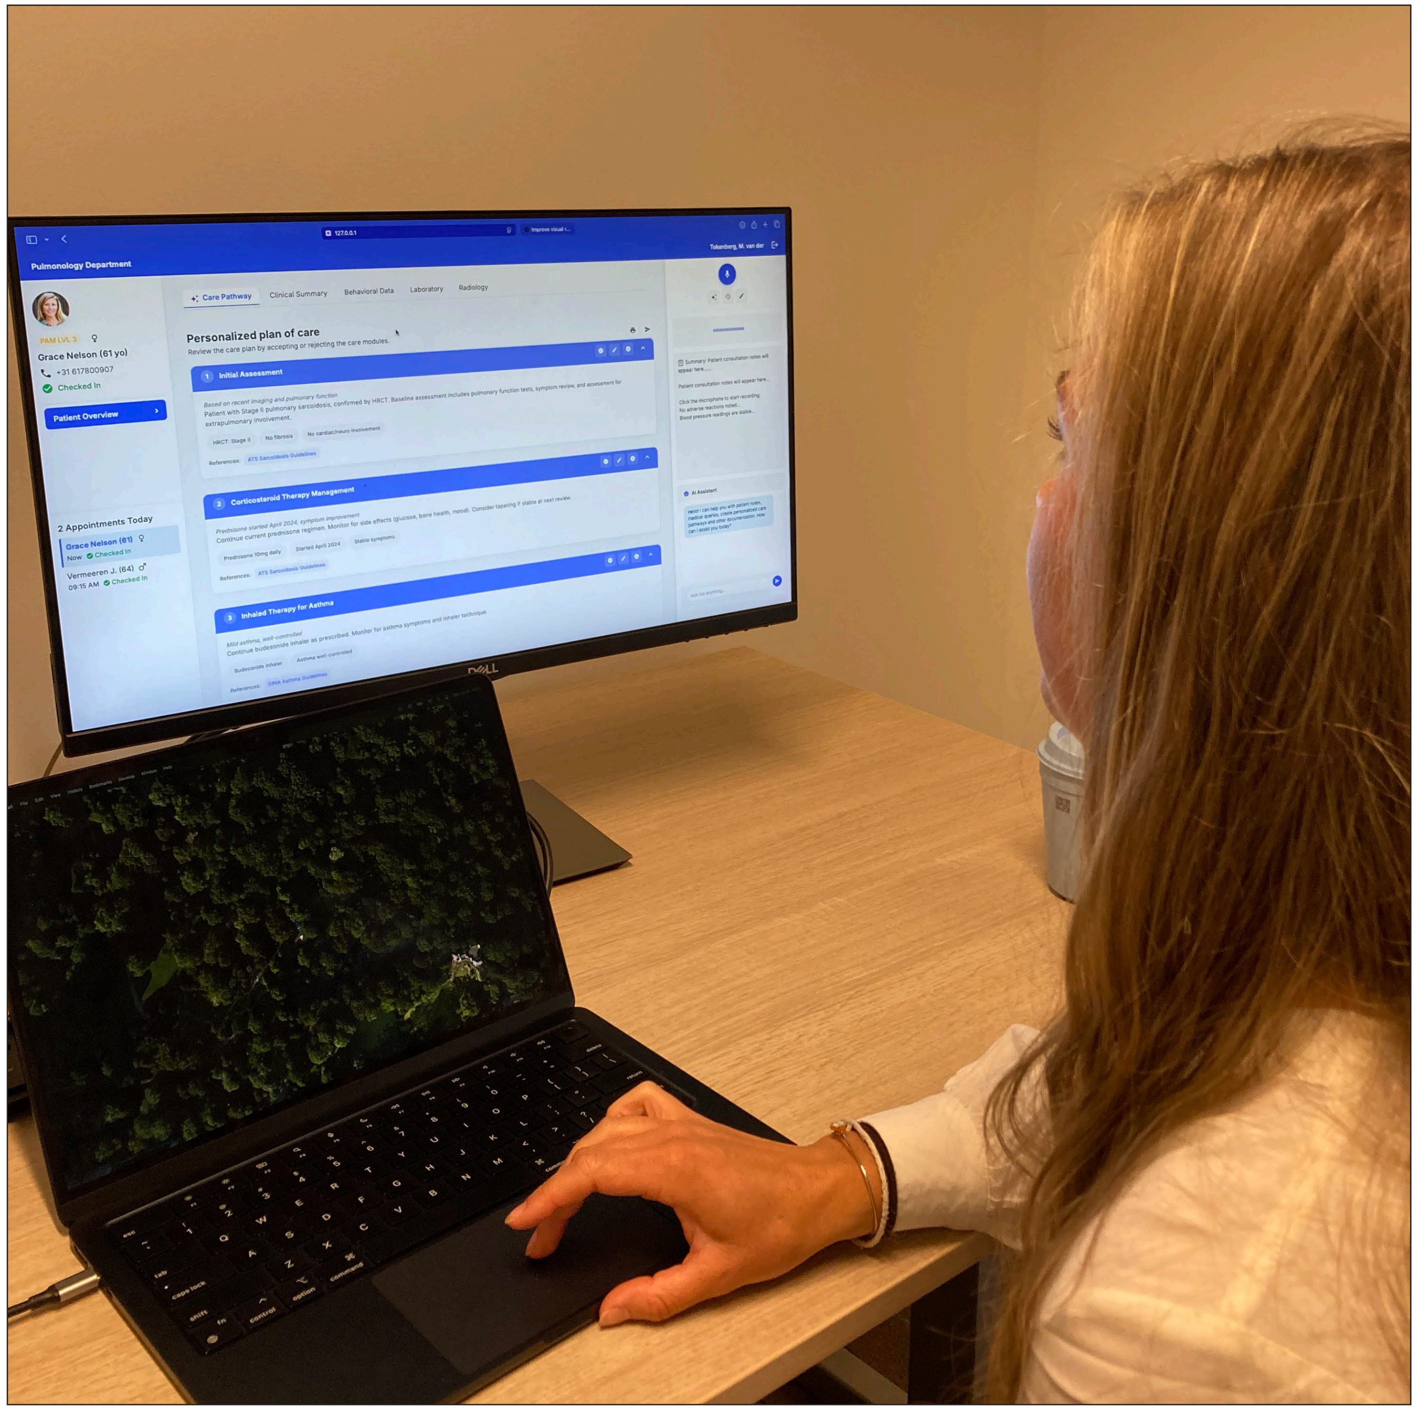

Fig 9.3 Clinician exploring the agent led personalization feature of the dashboard

> “You need to scroll through lots of data to read and understand what other professionals have said. And you’re repeating the same info over and over”

- All participants

Even role-based filters in HIX were seen as insufficient. The experts appreciated how this prototype was able to intelligently surface critical information in the form of clinical summaries in chronological order

> “Even with filters like surgical oncology, there’s still a lot of noise. This is more useful”

- PhD researcher in surgical oncology

Multidisciplinary ownership of the care modules

The clinicians emphasized that certain modules, especially those involving nutrition, physical activity, or psychological wellbeing should not be defaulted to physicians. They supported more dynamic role allocation across the care team.

> “Dietary counselling would be very beneficial if a dietitian were responsible for reviewing these…same for behavioral health”
>
> *- PhD researcher in surgical oncology*

They emphasized that the tool should be able to assign or suggest responsibilities based on module type & provider role.

Vision 2040 roadmap feels achievable

All three participants expressed confidence in the feasibility of the long-term vision. They considered the direction realistic, particularly if aligned with broader national digital infrastructure such as DigiZorg. Experts saw value in interoperable design that allows the system to integrate with current platforms, improving data continuity and clinical adoption.

> “That would be very achievable in the next 10 to 15 years. Having a more general app for different care centers even GP offices is the future. Right now, we have too many things”
>
> *- PhD researcher in surgical oncology*

9.3 Results and design implications

The feedback sessions revealed critical insights about the practical limitations and design opportunities within behavior-aware clinical systems. While the clinicians expressed appreciation for many aspects of the dashboard, their concerns pointed to deeper structural and workflow-related challenges. This section reflects on how these findings refine the Vision 2040 strategy and the future development of AI-powered tools in care planning.

Behavioral data requires context & timing

The clinicians recognized the value of behavioral data. However, they expressed hesitation about relying on it without clinical context. They warned against treating behavioral scores as objective facts. These insights carry meaning only when interpreted through a clinician’s evolving relationship with the patient.

Designers must embed behavioral signals directly within the clinical narrative. Isolated visualizations or separate tabs weaken their usefulness. Instead, clinicians should encounter this information alongside diagnosis, treatment history, and current care tasks. Behavioral data should guide attention, not dictate decisions. This insight addresses the first research question. Behavioral data can enrich care pathways when it informs human judgment without demanding new cognitive overhead.

AI should guide but never prescribe

The AI-generated care modules received praise for their modularity and traceability. However, clinicians drew a clear line. They accepted AI as a tool for surfacing options not as a mechanism for delivering care decisions.

Designers must resist the temptation to over-automate. The interface should treat clinicians as decision-makers, not validators of machine logic. Even small shifts in how suggestions appear such as changing labels from “recommended” to “optional” can alter perceptions of control. This finding responds to the second research question. AI can support behavior-aware decision-making only when clinicians feel in charge of the process.

Asynchronous features risk shifting burden not eliminating it

The dashboard’s asynchronous goal review feature opened a conversation about time. Although clinicians found the idea meaningful, they raised concerns about review fatigue. They feared a system that passively accumulates tasks would eventually be ignored or resented. Designing for asynchronous workflows requires more than notifications.

The system must actively support planning and prioritization. Features like review scheduling, auto-sorting by urgency, or integration with existing appointment structures become essential. Without these, asynchronous review becomes invisible work. This concern highlights the gap between emotional support and operational feasibility. Tools that increase patient engagement must also protect clinical time.

Shared interfaces demand shared responsibility

Experts endorsed giving patients access to their behavioral data. But the question of whether patients would benefit from such information is still yet to be answered. Each patient archetype identified in the earlier stages may respond differently to such information. This creates new challenges. If patients receive insights without guidance, they may misinterpret scores or become anxious.

Clinicians would then face new responsibilities as interpreters. Therefore, designers must build shared views that serve both patients and professionals. This requires plain language, transparent logic, and clear boundaries between what the system suggests and what the clinician endorses. Without this, shared access increases confusion and accountability risk.

Integration must come first, not last

Experts repeatedly stressed that adoption depends on integration. They described existing platforms like HIX as messy but established. Any new system must reduce friction, not add to it. A standalone tool, no matter how intelligent, will not succeed without fitting into existing workflows. Designers must treat interoperability as a foundation, not an afterthought. Adopting standards like FHIR, aligning with DigiZorg infrastructure, and mimicking familiar layouts are strategic moves that determine clinical acceptance.

Key Takeaways

These findings shift the project’s focus from interface design to system embedding and real-world implementation. The dashboard can only succeed if it acts as a translator converting behavioral signals into clinical insight without disrupting routines. AI can only support care when it adapts to professional judgment. Vision 2040 is achievable only through alignment and not disruption.

By viewing behavioral insights as contextual signals and AI as a supportive layer, the design respects both the complexity of care and the autonomy of clinicians. This validates the promise of data-enabled design but only when it follows the pace and culture of existing clinical practice.

9.4 Reflections

I went into the interviews hoping to confirm the value of my work. I came out with a clearer understanding of its limits. Clinicians appreciated the direction, but they asked hard questions that forced me to rethink parts of the system. I learned that good intentions do not always translate into useful tools. When I introduced asynchronous goal review, I believed it would support empathy. Clinicians saw the risk of hidden workload. I had to accept that even helpful features can add friction when they lack structure or timing.

I also learned how important trust is when designing with AI. The experts welcomed AI suggestions only when they could trace the logic and control the outcome. They refused to accept decisions made on their behalf. I realized I needed to make clinical judgment more visible, not less. This phase taught me to treat feedback as a design tool. I built the prototype to spark reaction, not just approval. I asked questions and let the clinicians lead the critique. Their input helped me shift from designing features to designing workflows. That shift changed how I think about impact.

These conversations grounded my vision. They showed me that clinical change does not come from new technology alone. It comes from careful alignment with habits, values, and systems. I now see personalization not as automation, but as a way to make decisions more human.

9.5 Conclusions

This chapter began with two guiding research questions. What is the impact of providing structured behavioral data alongside clinical patient information? How do clinicians view AI-assisted, behavior-aware decision support tools in practice?

Through expert walkthroughs with healthcare professionals at Erasmus MC, I explored these questions using a working prototype of the Vision 2040 clinician dashboard. These sessions tested how behavioral data, intelligent care planning, and shared decision-making could fit into real clinical workflows. Rather than gathering surface-level reactions, this chapter focused on uncovering the deeper tensions between opportunity and adoption.

One of the clearest findings was that clinicians recognize the value of behavioral insights when these are presented with context and clarity. They want to know how patients feel, think, and act, but only when that information arrives at the right moment and in the right format. Behavioral scores became meaningful when they complemented clinical experience rather than competing with it.

The AI-driven care pathway generator pushed this insight further. Experts responded positively to the idea of receiving suggestions grounded in behavior and lifestyle. They appreciated the link to clinical guidelines and the modular structure of the output. The interface gained trust because it supported judgment rather than replacing it.

The feedback sessions also challenged some assumptions. Features like asynchronous goal review introduced unintended concerns about time and responsibility. What I saw as emotional support, clinicians saw as potential overload. This pushed the design toward more structured review systems that respect the limits of daily workflow. Personalization only works when it does not fragment attention or delay action.

Finally, the experts placed strong emphasis on integration. They viewed the dashboard as promising because it felt calm, focused, and familiar. They reminded me that innovation in healthcare does not begin with novelty. It begins with alignment in language, habits, systems, and professional roles.

This chapter validated the core ideas behind the Vision 2040 roadmap. It showed that clinicians want personalization, as long as it fits their context. They welcome AI, but only when it strengthens their agency. They value behavioral data, but only when it is timely, structured, and clinically relevant.

In the next chapter, I reflect on the larger journey of this thesis. I discuss the limitations of my approach and suggest ways forward for implementing behavior-aware systems in everyday care.

Chapter 10 Conclusions

This chapter brings together the answers to the central research question of this thesis. It explores how data enabled design can support the creation of personalized care pathways for chronic patients and reduce the burden on healthcare providers.

In the previous chapters, I investigated this question by working with real patient narratives, building new design methods, and testing prototypes with clinicians. I developed a pipeline that turned online stories into structured insights. I built a dashboard that showed how these insights could support clinical decisions. I engaged with clinicians to test the relevance and reliability of the work. Throughout the process, I addressed concerns around data quality, privacy, ethics, and bias. These challenges shaped the design decisions and guided the scope of what this thesis could deliver.

In this chapter, I begin by explaining the three core pillars that supported the design approach. I then describe the limitations that emerged during the research. I follow this with targeted recommendations for academic researchers and hospital leaders who want to build on this work. I end with a personal reflection on the journey, the motivation that sustained it, and the lessons I carry forward.

10.1 Discussion

This study showed that data enabled design can create behavior aware care pathways using real patient narratives while reducing clinician burden. The work introduced three practical pillars that made this approach possible. Each pillar addressed a different part of the research question and together they formed a complete and grounded design strategy.

Iterative data verification

The first pillar focused on data quality. The persona pipeline used structured prompts and a chain of verification loop to test every generated detail against original patient posts. This loop helped reduce hallucinations and uncovered gaps that needed manual review. Chapter 4 described how the model resampled posts until the output met accuracy criteria without ever exposing identifiable information. This process kept the generated personas contextually grounded, accurate, and safe to use. The system also met key privacy standards by working with de-identified data inside a secure environment.

Clinician co-creation

The second pillar addressed ethical oversight. The co-creation sessions with pulmonologists and oncologists tested whether the model's behavioral insights matched clinical reality. These sessions confirmed many patterns but also revealed blind spots. For example, clinicians flagged that passive patients or those who lack digital access were missing from early versions. They also emphasized that behavior changes over time and that static profiles do not reflect real care journeys. Chapter 5 showed how their input shaped the final archetypes and made them usable in real care planning. The sessions also strengthened trust and ensured that domain expertise remained central.

Adaptive interface design

The third pillar focused on usability and integration. The patient-facing tools collected mood and goal setting data in a way that felt natural and non-intrusive. On the clinician side, the dashboard plugin for the HIX system displayed behavioral trends next to medical records and allowed AI-generated care suggestions. Chapter 9 documented how clinicians responded positively to this design, especially because the system made all suggestions transparent, editable, and traceable. These features reduced cognitive load, preserved clinical autonomy, and aligned with current regulatory expectations.

Synthesis

Together, these pillars showed that careful verification can improve data reliability, clinician engagement can expose bias and improve fairness and interface design can support clinical workflows without disrupting them. The result is a personalized care model that respects patient diversity, supports ethical decision-making, and reduces unnecessary work for healthcare providers. This approach provides a grounded and actionable answer to the central research question.

10.2 Limitations

This study strengthened the understanding of data enabled design in chronic care. However, several limitations affect how far these findings can travel across different patient groups and clinical settings.

Data Representativeness
The dataset included over 13,000 blog posts from three online peer support forums. This scale added depth and variety to the behavioral signals, but it focused on people who already engage with digital communities. It excluded those who avoid online spaces, struggle with digital literacy, or face language barriers. As a result, the patterns observed may not reflect the experiences of older adults, migrants, low-income groups, or patients with limited digital access. This gap reduces equity as the model scales to broader populations.

Generalizability of behavior profiles
The behavioral profiles in this thesis came from patients with chronic respiratory and cancer conditions. These pathways follow relatively stable and predictable patterns over time. The same approach may not apply to patients with sudden acute events, complex multimorbidity, or pediatric conditions where behaviors shift rapidly. Cultural differences also limit generalization. Health beliefs, family roles, and reimbursement structures change across countries. To validate the profiles further, future work needs to test them across more diverse clinical and social contexts.

Technical constraints of large language models
The persona pipeline relied on OpenAI's GPT4o large language model. Early runs produced hallucinations and generic content, which led to the development of a chain of verification loop. This step improved accuracy but required more compute time and constant manual review. Some long forum threads also exceeded token limits, which forced content to be summarized and caused the loss of emotional nuance.

Moreover, there is always the issue of AI models being biased since we did not train the model in identifying behavioral characteristics. Though gaurdrails were used in the form of prompts (for e.g. defining exactly what behavioral characteristics should be extracted) for analysis, further research should focus on unbiased ways of training the LLM model before using it for behavioral analysis.

Organizational readiness
The clinician dashboard ran in a sandbox using synthetic patient records. It was not deployed in a live clinical environment. Full implementation will require secure integration with hospital systems, training for staff, and ongoing technical support. Clinicians who tested the dashboard gave positive feedback, but they also flagged real concerns about workload, alert fatigue, and the time needed to adopt new tools. Without investment in change management and leadership support, the system may not move beyond the prototype stage.

10.3 Implementation and future Recommendations

The findings from this thesis offer practical guidance for two key audiences: academic researchers seeking to build on the methods, and hospital leadership looking to bring behavior-aware tools into care settings. Rather than propose immediate commercial deployment, these recommendations focus on how to enhance the evidence base, support clinician trust, and integrate ethically grounded design within existing healthcare systems.

10.3.1 Academic researchers

Refine the role of online patient narratives in clinical design
This thesis showed that peer support forums hold rich, behaviorally relevant signals across thousands of patients. Yet clinical design often ignores them. Future research should examine how to systematically integrate such narratives alongside existing PROMs and structured questionnaires. One direction could involve blending forum-derived personas with clinician-rated behavioral assessments to explore alignment and divergence over time.

Investigate manual verification as a design research method
Manual verification was not a bottleneck but a design tool in this project. Researchers should explore how human-in-the-loop validation of LLM outputs can surface tacit criteria, ethical red flags, or overlooked biases in patient representations. In this thesis, verification not only improved output quality but also clarified what "accuracy" meant in a clinical communication context.

Develop frameworks for behavior aware prototyping
The patient-facing flows and clinician dashboards did not present final products, but research probes that helped test how behavior insights travel across interfaces. Academic design teams should now formalize these flows into reusable frameworks: how to embed adaptive goal setting, when to visualize sentiment shifts, and where to show AI-generated suggestions without undermining clinician trust. These questions need deeper study in both HCI and health informatics domains.

Explore how behavioral data shapes clinical interpretation
Expert feedback sessions revealed that clinicians interpret behavioral cues not as final answers but as conversation starters. Future studies should explore how behavior flags affect diagnostic reasoning, shared decision-making, or triage in real practice. These investigations can help determine whether integrating such data improves care continuity, empathy, or efficiency and for whom.

10.3.2 Hospital leadership

Support sandboxed experimentation inside existing EHR environments
This work showed that AI-generated behavioral insights gain clinician trust when embedded in familiar systems. The clinician dashboard plugin was designed to fit into HIX, and walkthroughs confirmed its relevance in real case workflows. Hospitals should create protected sandbox spaces where innovations like this can run in read-only or training modes, allowing safe iteration without disrupting live care.

Treat AI transparency and clinician autonomy as core design principles
Across the thesis, clinicians consistently responded best when suggestions came with context, not conclusions. The dashboard clearly separated patient input from AI interpretation, kept all suggestions editable, and avoided automation bias. Any future implementation should preserve this approach and ensure clinicians remain in full control. These practices also align closely with the EU AI Act's transparency requirements for high-risk systems.

Build local capacity for human-in-the-loop data validation
Behavioral insights require context. The project's verification loop relied not only on AI feedback but also on judgment from a design researcher. Hospitals exploring similar solutions should invest in teams that combine clinical, design, and data expertise to monitor and improve these systems locally. This is not just a technical requirement but a governance safeguard.

Focus on behavioral data that improves relational care
Clinicians at Erasmus MC noted that the most valuable behavioral signals were not predictive, but relational. They helped explain missed appointments, hesitations about treatment, or changing attitudes across a care journey. Hospitals should prioritize tools that make these signals visible and actionable tools that enhance empathy, not efficiency alone.

These recommendations build on validated prototypes, real clinician feedback, and extensive qualitative synthesis. They do not assume that AI or data-driven personalization will replace clinical expertise. Instead, they argue that with careful design, transparent workflows, and institutional commitment, behavioral insights can support more human, more relational, and more effective care especially for patients with complex chronic conditions.

10.4 Personal reflections on the overall journey

This thesis began with a simple goal to explore how data enabled design could help patients and reduce pressure on healthcare systems. I had no idea how much this journey would change the way I think, design, and approach life itself. The biggest surprise came early, when I started exploring real patient blogs. I expected scattered thoughts and disconnected stories. What I found instead was nuance, honesty, and rich context with thousands of patients, each writing from a place of struggle, hope, and resilience. Reading these blogs felt like listening in on something intimate. When I realized that AI could analyze and make sense of this emotional complexity, I felt something change. It became clear to me that these were not just stories. They were signals, and if handled carefully, they could become the foundation for personalized, humane, and responsive healthcare.

But getting from insight to solution was not easy. One of the hardest moments came when I had to translate all the behavioral, systemic, and contextual knowledge into a concrete design. I spent weeks trying to imagine the right future use case. I wanted to build something meaningful, something rooted in technology but driven by care. I got stuck. But slowly, my systems research helped me frame the problem more clearly. That clarity gave me the energy to design forward, and eventually, the clinician dashboard came into view. That moment of transition from uncertainty to direction remains one of the most fulfilling phases of this work.

I also grew technically and intellectually. I coded the data pipeline myself. I worked through verification processes, reviewed outputs line by line, and spent long nights and multiple cups of coffee debugging, retraining, and rewriting prompts. What seemed impossible at the start became my everyday routine. More than anything, this work taught me that designers are not separate from data or systems. We belong right inside them and have the knowledge necessary to shape how they behave.

This work also changed how I think about AI and privacy. As I watched the models build portraits of patients from just a few lines of text, I realized how much can be inferred from language alone (just one simple sentence was enough). It made me far more cautious about my own digital footprint. At the same time, it reinforced the idea that personalized care is not just a feature or a trend. It is a question of equity. Every person deserves to be understood in their own context. This thesis became my way of exploring how to design for that kind of dignity.

Some of the most joyful moments came during user testing. I remember showing a prototype of the dashboard to a clinician who immediately exclaimed, “Wow! This is so much nicer than HIX!” That one comment made me smile for days. It reminded me why I chose to design in the first place, to build concepts that bring clarity, ease, and maybe even joy into hard situations. But I also saw how difficult it is to engage clinicians, not because they lack interest, but because they are overburdened, understaffed, and boxed in by rigid systems. Getting meetings took weeks. Implementation will take years.

Even so, I was surprised by how open they were. When I finally presented my work at Erasmus MC, I expected skepticism. Instead, I found curiosity and enthusiasm. They wanted to explore this future. They just needed help navigating the slow-moving world of hospital regulations. This topic is personal for me. I have seen care delivery fail people I love in the form of misdiagnoses, wrong medications, and preventable suffering. My grandmother passed after being given the wrong drug. My grandfather was a victim of cancer. My younger brother struggled through his childhood due to poor clinician decision making. These stories are why I chose healthcare. They are why I will continue. If I had more time, I would refine the dashboard and test it with a wider group of clinicians. I would build it into a real product and take the next steps towards building my own startup. That vision is still alive.

I hope readers remember not just the methods or results in this thesis, but the mindset behind it. The belief that data can carry empathy. That AI can be transparent. That care can be personalized without losing humanity. That design, when done right, can shape futures worth living in.

This thesis was never just about technology.
It was about what I chose to do with it.

References

1. 36-Item Short Form Survey Instrument (SF-36). (n.d.). RAND. https://www.rand.org/health-care/surveys_tools/mos/36-item-short-form/survey-instrument.html

2. Abdelraheem, O., Salama, M., & Chun, S. (2024). Impact of digital interventions and online health communities in patient activation: Systematic review and meta-analysis. International Journal of Medical Informatics, 188, 105481. https://doi.org/10.1016/j.ijmedinf.2024.105481

3. Anand, R. (2024). Enhancing Patient Care Pathways through AI-Driven Data Science and Program Management Strategies. International Journal of Media and Networks, 2(8), 01–06. https://doi.org/10.33140/ijmn.02.08.04

4. Artificial Intelligence in healthcare. (2025, June 10). Public Health. https://health.ec.europa.eu/ehealth-digital-health-and-care/artificial-intelligence-healthcare_en?utm_source=chatgpt.com#eu-legislation-shaping-ai-in-healthcare

5. Bodenheimer, T., Wagner, E. H., & Grumbach, K. (2002). Improving primary care for patients with chronic illness. JAMA, 288(15), 1909. https://doi.org/10.1001/jama.288.15.1909

6. Bour, S. S., Raaijmakers, L. H. A., Bischoff, E. W. M. A., Goossens, L. M. A., & Mölken, M. P. M. H. R. (2023). How Can a Bundled Payment Model Incentivize the Transition from Single-Disease Management to Person-Centred and Integrated Care for Chronic Diseases in the Netherlands? International Journal of Environmental Research and Public Health, 20(5), 3857. https://doi.org/10.3390/ijerph20053857

7. Bucher, A. (2023). The Patient Experience of the Future is Personalized: Using Technology to Scale an N of 1 Approach. Journal of Patient Experience, 10. https://doi.org/10.1177/23743735231167975

8. Büssing, A., Balzat, H., & Heusser, P. (2010). Spiritual needs of patients with chronic pain diseases and cancer - validation of the spiritual needs questionnaire. European Journal of Medical Research, 15(6), 266. https://doi.org/10.1186/2047-783x-15-6-266

9. Carroll, Á., Collins, C., & McKenzie, J. (2025). Physician wellbeing in a national rehabilitation hospital, a qualitative study utilizing Maslow's hierarchy of needs as a framework for analysis. BMC Health Services Research, 25(1). https://doi.org/10.1186/s12913-025-12310-x

10. Chen, S., Guevara, M., Moningi, S., Hoebers, F., Elhalawani, H., Kann, B. H., Chipidza, F. E., Leeman, J., Aerts, H. J. W. L., Miller, T., Savova, G. K., Gallifant, J., Celi, L. A., Mak, R. H., Lustberg, M., Afshar, M., & Bitterman, D. S. (2024). The effect of using a large language model to respond to patient messages. The Lancet Digital Health, 6(6), e379–e381. https://doi.org/10.1016/s2589-7500(24)00060-8

11. Chen, Y., Lehmann, C. U., & Malin, B. (2024). Digital information ecosystems in modern care coordination and patient care pathways and the challenges and opportunities for AI solutions. Journal of Medical Internet Research, 26, e60258. https://doi.org/10.2196/60258

12. De Filippis, G. M., Rinaldi, A. M., Russo, C., & Tommasino, C. (2024). Advanced Topic Modeling in Genomics: Towards Personalized Dietary Recommendations through BERTopic Analysis. In Lecture notes in computer science (pp. 3–17). https://doi.org/10.1007/978-3-031-78093-6_1

13. Dhuliawala, S., Komeili, M., Xu, J., Raileanu, R., Li, X., Celikyilmaz, A., & Weston, J. (2024). Chain-of-Verification reduces hallucination in large language models. Findings of the Association for Computational Linguistics: ACL 2022, 3563–3578. https://doi.org/10.18653/v1/2024.findings-acl.212

14. Džakula, A., Relić, D., & Michelutti, P. (2022). Health workforce shortage – doing the right things or doing things right? Croatian Medical Journal, 63(2), 107–109. https://doi.org/10.3325/cmj.2022.63.107

15. European Health Data Space Regulation (EHDS). (2025, June 10). Public Health. https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en

16. Foreman, K. J., Marquez, N., Dolgert, A., Fukutaki, K., Fullman, N., McGaughey, M., Pletcher, M. A., Smith, A. E., Tang, K., Yuan, C.-W., Brown, J. C., Friedman, J., He, J., Heuton, K. R., Holmberg, M., Patel, D. J., Reidy, P., Carter, A., Cercy, K., . . . National Institute of Public Health, Cuernavaca, Mexico. (n.d.). Forecasting life expectancy, years of life lost, and all-cause and cause-specific mortality for 250 causes of death: reference and alternative scenarios for 2016–40 for 195 countries and territories. The Lancet, 392, 2052–2090.

17. Hand, C. (2016, July 1). Measuring health-related quality of life in adults with chronic conditions in primary care settings: Critical review of concepts and 3 tools. https://pmc.ncbi.nlm.nih.gov/articles/PMC4955103/

18. Health at a Glance 2023. (2023). In Health at a glance. https://doi.org/10.1787/7a7afb35-en

19. Hjelm, J. (2010). The dimensions of health: Conceptual Models. Jones & Bartlett Learning.

20. Iglesias, F. H., Celada, C. A., Navarro, C. B., Morales, L. P., Visus, N. A., Valverde, C. C., Dorado, V. M., Muñoz, M. L. M., Pueyo, C. B., Ramírez, E. L., Villaroel, R. M., Jordana, N. M., & Simó, J. M. B. (2018). Complex care needs in multiple chronic conditions: population prevalence and characterization in primary care. a study protocol. International Journal of Integrated Care, 18(2). https://doi.org/10.5334/ijic.3292

21. Javaid, M., Haleem, A., & Singh, R. P. (2024). Health informatics to enhance the healthcare industry's culture: An extensive analysis of its features, contributions, applications and limitations. Informatics and Health, 1(2), 123–148. https://doi.org/10.1016/j.infoh.2024.05.001

22. Jung, J. (2023). Developing data-enabled design in the field of digital health. repository.tudelft.nl. https://doi.org/10.4233/28c38358-a1ca-423c-97fb-841471138e56

23. Karimi, M., Tsiachristas, A., Looman, W., Stokes, J., Van Galen, M., & Mölken, M. R. (2021). Bundled payments for chronic diseases increased health care expenditure in the Netherlands, especially for multimorbid patients. Health Policy, 125(6), 751–759. https://doi.org/10.1016/j.healthpol.2021.04.004

24. Khalifa, M., & Alswailem, O. (2015). Clinical Pathways: Identifying development, implementation and evaluation challenges. Studies in Health Technology and Informatics. https://doi.org/10.3233/978-1-61499-538-8-131

25. Korade, N. B., Salunke, M. B., Bhosle, A. A., Kumbharkar, P. B., Asalkar, G. G., & Khedkar, R. G. (2024). Strengthening sentence similarity identification through OpenAI embeddings and deep learning. International Journal of Advanced Computer Science and Applications, 15(4). https://doi.org/10.14569/ijacsa.2024.0150485

26. Kwan, J. (2007). Care pathways for acute stroke care and stroke rehabilitation: From theory to evidence. Journal of Clinical Neuroscience, 14(3), 189–200. https://doi.org/10.1016/j.jocn.2006.01.026

27. Lidwell, W., Holden, K., Butler, J., Butler, J. C., & Lidwell, L. W. (n.d.). Universal principles of design (pp. 8–268).

28. López-Salas, M., Yanes-Roldán, A., Fernández, A., Marín, A., Martínez, A. I., Monroy, A., Navarro, J. M., Pino, M., Gómez, R., Rodríguez, S., Garrido, S., Cousillas, S., Navas, T., Lapeña, V., & Fernández, B. (2024). End-of-life care needs in cancer patients: a qualitative study of patient and family experiences. BMC Palliative Care, 23(1). https://doi.org/10.1186/s12904-024-01489-1

29. Medical Device Coordination Group. (2021). Guidance on classification of medical devices. In Medical Device Coordination Group Document MDCG 2021- 24. https://health.ec.europa.eu/system/files/2021-10/mdcg_2021-24_en_0.pdf

30. Mullick, A., Martin, J., & Sallnow, L. (2013). An introduction to advance care planning in practice. BMJ, 347(oct21 3), f6064. https://doi.org/10.1136/bmj.f6064

31. Nedungadi, P., Veena, G., Tang, K., Menon, R. R., & Raman, R. (2025). AI techniques and applications for online social networks and media: insights from BERTopic modeling. IEEE Access, 1. https://doi.org/10.1109/access.2025.3543795

32. Po, H., Chu, Y., Tsai, H., Lin, C., Chen, C., & Huei-Ming, M., MA. (2024). Efficacy of remote health monitoring in reducing hospital readmissions among High-Risk Postdischarge patients: Prospective cohort study. JMIR Formative Research, 8, e53455. https://doi.org/10.2196/53455

33. Regulation - EU - 2024/1689 - EN - EUR-LEX. (n.d.). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32024R1689

34. Roesel, I., Froehlich, D., Joos, S., Valentini, J., Mauch, H., & Martus, P. (2024). The Patient Activation Measure-13 (PAM-13) in an oncology patient population: psychometric properties and dimensionality evaluation. Health and Quality of Life Outcomes, 22(1). https://doi.org/10.1186/s12955-024-02255-w

35. Sahoo, P., Singh, A. K., Saha, S., Jain, V., Samrat Mondal, & Aman Chadha. (2021). A Systematic survey of prompt engineering in large language Models: Techniques and applications. In Department of Computer Science and Engineering, Indian Institute of Technology Patna [Journal-article].

36. Sikka, R., Morath, J. M., & Leape, L. (2015). The Quadruple Aim: care, health, cost and meaning in work. BMJ Quality & Safety, 24(10), 608–610. https://doi.org/10.1136/bmjqs-2015-004160

37. SolvEdge. (n.d.). Healthcare in 2040 – the future of the hospital. https://www.solvedge.com/Healthcare-in-2040

38. Sushil, M., Suster, S., & Daelemans, W. (2021). Are we there yet? Exploring clinical domain knowledge of BERT models. ACL Anthrology. https://doi.org/10.18653/v1/2021.bionlp-1.5

39. Vollset, S. E., Ababneh, H. S., Abate, Y. H., Abbafati, C., Abbasgholizadeh, R., Abbasian, M., Abbastabar, H., Magied, A. H. a. a. A., ElHafeez, S. A., Abdelkader, A., Abdelmasseh, M., Abd-Elsalam, S., Abdi, P., Abdollahi, M., Abdoun, M., Abdullahi, A., Abebe, M., Abiodun, O., Aboagye, R. G., . . . Murray, C. J. L. (2024). Burden of disease scenarios for 204 countries and territories, 2022–2050: a forecasting analysis for the Global Burden of Disease Study 2021. The Lancet, 403(10440), 2204–2256. https://doi.org/10.1016/s0140-6736(24)00685-8

40. Voogt, E., Van Leeuwen, A. F., Visser, A. P., Van Der Heide, A., & Van Der Maas, P. J. (2005). Information needs of patients with incurable cancer. Supportive Care in Cancer, 13(11), 943–948. https://doi.org/10.1007/s00520-005-0823-6

41. Wang, T., Molassiotis, A., Chung, B. P. M., & Tan, J. (2018). Unmet care needs of advanced cancer patients and their informal caregivers: a systematic review. BMC Palliative Care, 17(1). https://doi.org/10.1186/s12904-018-0346-9

42. WHOQOL-BREF. (n.d.). https://www.who.int/tools/whoqol/whoqol-bref

43. Winter, V., Schreyögg, J., & Thiel, A. (2020). Hospital staff shortages: Environmental and organizational determinants and implications for patient satisfaction. Health Policy, 124(4), 380–388. https://doi.org/10.1016/j.healthpol.2020.01.001

44. World Health Organization: WHO. (2024, December 23). Noncommunicable diseases. https://www.who.int/news-room/fact-sheets/detail/noncommunicable-diseases

45. Yablonski, J. (n.d.). Jakob's Law | Laws of UX. Laws of UX. https://lawsofux.com/jakobs-law/

Designing the future of healthcare

AI x
data x
human touch.