Deliverable D 4.2

Co-Development of ICT Solutions

Authors: Maria Nefeli Kousta (ICCS)

Nikos Paraskevopoulos (ICCS)

Ioannis Katsipis (ICCS)

Hlias Doukas (ICCS)

https://www.dietwise.eu

Funded by the European Union

This project is funded by the European Union's *Horizon 2020 program* under grant agreement No. 101181692

This work is dedicated to the memory of Justina Baršytė, author of the DietWise project idea, whose vision and commitment were invaluable to this project.

Project information

Program:	Horizon Europe
Topic:	HORIZON-CL6-2024-FARM2FORK-01-5
Type of action:	HORIZON-RIA HORIZON Research and Innovation Actions
Grant Agreement #:	101181692
Project title:	Systemic Solutions to Enhance Healthy and Sustainable Food Provision and Cooking at Home
Project Name:	DietWise
Project Start Date:	2024-11-01
Project End Date:	2027-10-31

Document information

Document name:	Co-Development of ICT Solutions
Related Work Package:	WP 4
Related Task:	Task 4.2
Related Deliverables:	D 4.2
Author(s):	Maria Nefeli Kousta, Nikos Paraskevopoulos, Ioannis Katsipis, Hlias Doukas
Reviewer(s):	Ebo Botchway (KUL), Zivile Kaminskiene (AdC)
Submission date:	2026-02-27
Dissemination level:	Public

Document history

Version	Date	Changes	Responsible partner
v0.1	2025-12-10	TOC	ICCS
V0.2	2026-01-09	Draft ready for internal review	ICCS, SAFE, VIGL, IHU, PROL, PHB, VU
V0.3	2026-02-20	Review – Comments and Suggestion	KUL, AdC
V0.4	2026-02-27	Final Version ready for submission	ICCS, SAFE, VIGL, IHU, PROL, PHB, VU

Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

Abbreviations

Abbreviation	Full Form
AI	Artificial Intelligence
LLM	Large Language Model
RCA	Responsible Cooking Alliance
RW	Recipe Watch
ICT	Information and Communication Technology
FBDGs	Food-based dietary guidelines
GBD	Global Burden of Disease
WP	Work Package
RAG	Retrieval-Augmented Generation
PHB	Visiomenes Sveikatos Biuras
VIGL	Vlaams Institut voor Gezondheitdspromotie
IHU	Diethnes Panepistimio Ellados
SAFE	Safe Food Advocacy
PROL	Astiki Mi Kerdoskopiki Etaireia Proliptikis Perivallontikis kai Ergasiakis Iatrikis
VU	Vilniaus Universitetas
GBD	Global Burden of Disease

Executive Summary

This deliverable presents the development of ICT solutions, as they were shaped through extensive cooperation with stakeholders across all pilot countries, as well as the nutritional experts of the consortium. It outlines the framework of co-development activities, namely their first main stage: *co-development for customization,* as well as the implementation of the workshops in each pilot country and the key findings and insights emerged. It provides all the outcomes of the extensive discussions with the nutritional experts that played a central role in co-developing and adapting the ICT solutions, highlighting the alignment with national guidelines, the extent of personalization and the scoring mechanism. The document provides an explicit description of the two applications -RecipeWatch (RW) and Responsible Cooking Alliance (RCA)- from their initial conceptualization to their core architecture. It also describes the framework of their usability testing procedure.

Significant progress has been achieved in transforming end-user requirements into functional and meaningful features. After having conceptualized an initial version of the ICT solutions, stakeholders were presented the application mockups and provided feedback, which in turn was incorporated to the ICT design, ensuring that the applications respond effectively to their needs and expectations. The framework for co-development activities was standardized by SAFE -for influencers and citizens- and PROL -for vulnerable citizens, ensuring that all three population sub-groups are being acknowledged and appropriately addressed. The three pilot countries, namely Belgium, Greece and Lithuania implemented the workshops following the framework, documented outputs using a structured approach and captured identified needs, barriers, opportunities, and implications for tool design.

Future efforts will focus on further development and delivery of the ICT solutions in the context of WP5 (T5.2) and continuous feature refinement according to stakeholder feedback. With a solid foundation now in place, further refinements will be guided by ongoing stakeholder feedback and carried out to the extent feasible. While adjustments may still be incorporated, the current stage has established a comprehensive and robust baseline that already addresses the identified stakeholders’ needs.

Table of Contents

List of Figures

List of Tables

1. INTRODUCTION

This deliverable (D4.2) reports the consolidated outcomes of WP4 related to the co-development of the DietWise ICT solutions. It documents how stakeholder insights and expert input are transformed into meaningful features, and how these features are reflected in the initial design and technical implementation of the two DietWise user-facing applications: Recipe Watch (RW) for citizens and Responsible Cooking Alliance (RCA) for influencers. D4.2 directly contributes to project objective O3 (Digital social innovations and AI-based apps to empower citizens) and supports KPIs 16-19 by establishing the foundations required to achieve them by the delivery of D5.2.

D4.2 is positioned as a bridge between concept co-creation and technical realization. On the one hand, it operationalizes the needs, barriers, and opportunities captured through early project activities (notably WP2, WP3, and T4.1) into concrete application workflows and feature sets. On the other hand, it provides the technical basis required for subsequent pilot planning and implementation activities (WP5 and WP7), by describing the current architecture, implementation approach, and the main technical restrictions and mitigation measures identified at this stage.

In line with the Grant Agreement requirements, this deliverable also includes the planned approach for usability testing of both RW and RCA. While the interactive prototypes and beta releases are not yet available at the time of writing, the deliverable defines the intended testing strategy and tooling to ensure that usability evaluation can be executed in a structured and auditable way once the relevant versions are ready.

1.1. PURPOSE AND SCOPE

The purpose of this deliverable is to document the first-stage development and technical grounding of the DietWise ICT solutions developed under WP4, and specifically Task T4.2 to support the initiative of the DietWise project in line with the needs of relevant stakeholders. **The scope of this deliverable covers the methodology for co-development and customization.** It describes the standardized approach used to engage stakeholders (citizens, influencers, and vulnerable citizens) in structured co-development activities, focusing on translating user needs into actionable functional requirements and interface expectations. It explains how stakeholder input is converted into application functionality, including the definition of key workflows (e.g., recipe submission/extraction, presentation of ingredient substitutions, user feedback/acceptance mechanisms, transparency indicators, and influencer-facing compliance feedback). It outlines the outcomes of the extensive discussions with the nutritional experts that played a central role in co-developing and adapting the ICT solutions, highlighting the alignment with national guidelines, the extent of personalization and the scoring mechanism. Additionally, it provides an overview of the current technical architecture of the DietWise framework supporting RW and RCA, including the common backend infrastructure, identity management and authentication approach, database separation principles, AI execution components, and recipe rendering/extraction mechanisms. In parallel, it gives a quality-risk perspective by identifying technical restrictions affecting both solutions at this stage (e.g., variability of recipe website structures, extraction robustness, and server-side rendering constraints) and outlines mitigation strategies consistent with the project's iterative development approach. Finally, it defines how usability testing will be conducted within the timeframe of T5.2, once prototypes and beta versions become available. This includes testing on prototypes via user testing platforms and beta testing via distribution channels. The proposed approach ensures the establishment of a clear path to evidence-based usability improvements.

This deliverable is not intended to serve as a final specification or final evaluation report. Instead, it captures the baseline design and development status, providing traceability from stakeholder input to implementation decisions. In doing so, it establishes a solid foundation for iterative refinement of the solutions up to M18, when the beta versions are scheduled for release in accordance with the Grant Agreement (GA).

1.2. INTERCONNECTION WITH OTHER WORK PACKAGES AND DELIVERABLES

D4.2 is tightly connected to multiple work packages and acts as a synthesis point where upstream evidence is translated into ICT design choices and implementation decisions. WP2 generates the empirical foundation for ICT development, increasing understanding of the current state of behavioral aspects related to citizen needs, barriers, and drivers for uptake of nutrition guidelines, and enhancement of healthy and sustainable food provision. Its outputs are documented in D2.1-2.4. D4.2 builds on these insights by translating them into concrete design decisions for RW and RCA (e.g., supported workflows, trust and transparency features, onboarding requirements, and feedback mechanisms). WP3 develops and tests behavioral interventions to support adoption and sustained engagement (e.g., social influence, norms, cultural moderators, and accessibility for vulnerable groups). The outcomes of WP3 reported in D3.1 inform both the ICT feature roadmap and the presentation strategies in the user interfaces (e.g., how recommendations are framed, how transparency and norms may be operationalized, and how reported information can reduce resistance rather than trigger it). D4.2 supports this integration by documenting the feature-level implementation context in which these interventions can later be embedded and validated.

Within WP4, D4.2 is primarily associated with task T4.1, which provides initial stakeholder-led framing of needs and opportunities and T4.2 that defines and documents the co-development methodology and its early outcomes described in D4.1. Additionally, D4.2 is linked to T4.4 that integrates outputs from WP2–WP4 into a holistic framework (D4.4). D4.2 contributes the ICT-specific component of this integration by documenting the solutions’ structure, functionality, and technical constraints at this stage. WP5 operationalizes the DietWise framework into pilot implementation plans and prepares ICT solutions for pilot launch (WP7). D4.2 supports WP5 by providing the current state of ICT design and architecture, an initial risk identification relevant to pilot feasibility, and a usability testing plan that will produce actionable improvement cycles. In this sense, D4.2 is a key input to ensuring that the ICT solutions are not only conceptually aligned with stakeholder needs but also technically ready and iteratively refined for piloting.

1.3. METHODOLOGY AND STRUCTURE OF THE DELIVERABLE

This deliverable is structured to provide a traceable narrative from stakeholder-derived requirements to implemented ICT design and technical realization. The methodology is defined as follows:

- Section 2 describes how end-user requirements are transformed into meaningful features, including the framework for co-development activities and pilot-specific outcomes for each stakeholder group.
- Section 3 presents expert-driven inputs that guide adaptation, including alignment with national nutritional guidelines, definition of the level of personalization, and scoring mechanism for the recipes.
- Section 4 provides a detailed description of the two DietWise applications (RW and RCA) and the underlying architecture of the ICT solutions, including tools and technologies, backend and AI components, and the technical flow of suggestions. It also identifies key technical restrictions and mitigation directions.
- Section 5 outlines the usability testing strategy, covering prototype-based testing and beta testing for coded applications, establishing the evaluation approach that will be executed once the relevant versions become available (M18).

This structure ensures that the deliverable is both methodologically grounded (stakeholder and expert input) and implementation-oriented (architecture, restrictions, and validation plan), supporting the project’s progression toward pilot preparation and deployment.

2. END-USER REQUIREMENTS TRANSFORMED INTO MEANINGFUL FEATURES

To ensure that the ICT solutions effectively respond to real user needs, expectations, and everyday practices, a structured co-development process was implemented across all pilot countries. This process brought together key stakeholder groups, namely citizens, vulnerable citizens and influencers, and actively involved them in the identification and refinement of end-user requirements as obtained within the workshops. Through these activities, user insights were systematically translated into meaningful inputs for the design and customization of the project’s digital tools. This section presents the co-development framework developed by Safe Food Advocacy (SAFE) and Astiki Mi Kerdoskopiki Etaireia Proliptikis Perivallontikis kai Ergasiakis Iatrikis (PROL), its practical implementation by the pilot partners, and the main insights generated through the workshops. The findings are presented per stakeholder group and per pilot country, highlighting both shared patterns and context-specific needs. Together, these results form the basis for transforming identified requirements into concrete features, design principles, and functional priorities for ICT solutions.

2.1. FRAMEWORK OF CO-DEVELOPMENT ACTIVITIES FOR CUSTOMIZATION

Co-development activities were implemented through a series of workshops conducted across all pilot cities. The targeted participants were grouped into three key stakeholder categories: citizens, vulnerable citizens and influencers. The overall workshop methodology was developed by SAFE –for citizens and influencers- and by PROL for vulnerable citizens. Pilot country teams Visuomenes Sveikatos Biuras (PHB), Vlaams Institut voor Gezondheidspromotie (VIGL) and Diethnes Panepistimio Ellados (IHU) organized and facilitated the workshops in Belgium, Greece and Lithuania, convening the three crucial stakeholder groups for the project on agreed dates and involving them in structured co-development activities. This country-specific approach reflects the different levels of experience and capacity across the pilots in engaging stakeholders and therefore required dedicated tracks in each location. Thus, the methodology was developed as a flexible framework to help field partners plan and deliver workshops that respond to local contexts and stakeholder needs.

For the Belgian workshops, VIGL actively engaged a diverse range of stakeholder groups to ensure a comprehensive and inclusive framework. Engagement efforts focused on three primary pillars: regular citizens to capture broad public sentiment, vulnerable populations -specifically individuals from migration backgrounds and those living in poverty- to ensure the framework addresses systemic barriers and equity, and key influencers to help amplify the project's reach and impact. During these initial sessions, preliminary end-user requirements were successfully collected and highlight the specific needs and challenges faced by these groups, providing a foundational dataset for the subsequent stages of framework creation.

Co-creation workshops in Greece were performed by IHU, engaging citizens, vulnerable citizens, and food influencers in line with the methodology - outlined by SAFE and PROL - in separate workshops. More specifically, one workshop was held for citizens, two for food influencers and two for vulnerable citizens.

In Lithuania, co-creation workshops were conducted by the Vilnius City Municipality PHB, engaging citizens, vulnerable citizens (seniors), and food influencers in accordance with the methodology outlined by SAFE and PROL. These groups participated in separate workshops: one workshop for citizens, three for food influencers, and one for vulnerable citizens.

The following subsections present the detailed methodology applied during the workshops for each stakeholder group. It should be noted that no RW or RCA prototypes were trialed or tested during this phase; prototype evaluation is foreseen in a second round of feedback, to be conducted once both ICT solutions are available in beta versions. Through these workshops we identified stakeholder needs, expectations, and constraints, in order to use this input to guide and inform the ongoing development of the ICT solutions.

2.1.1. Influencers

The influencer workshops were structured as co-development discussions. The focus was on surfacing influencers' needs, barriers, and opportunities in relation to the Responsible Cooking Alliance (RCA) and using this input to steer RCA design and development. The sessions explored concepts, expectations, and workable safeguards. No RCA prototype was trialed or tested in these workshops. **Food influencers were approached, including bloggers, dieticians, and chefs** who are active on social media platforms. Recruitment primarily built on relationships established during earlier interviews conducted under Task 2.3 and was complemented by additional outreach to ensure adequate representation across profiles. Given influencers' limited availability, workshops were intentionally designed as short and focused sessions conducted in small-group formats, between 1 and 3 participants at a time. Where needed, a sequential format was used to capture input across staggered attendance and consolidate insights. Each workshop followed a structured flow of four phases:

1. Context setting and introductions
2. Rapid exploration of needs, barriers, opportunities
3. RCA concept discussion, including what would make it attractive and what would trigger scepticism
4. Summary and next steps, including possible future involvement

To ensure consistency across pilots, RCA-related discussions were guided by explicit prompts. Participants were invited to reflect on what features or functions should be added, removed or simplified, to articulate their main concerns and to identify what types of safeguards, evidence, or validation would be required to strengthen trust. Where feasible, rapid prioritization was applied to clarify which issues mattered the most to participants. Throughout the process, partner organizations documented workshop outputs using a structured approach, capturing identified needs, barriers, opportunities, and implications for tool design. Particular attention was given to factors influencing adoption, including credibility, usability, integration burden, and potential reputational risks.

2.1.2. Citizens

The citizen workshops were designed as co-creation discussions, with the aim of actively involving participants in shaping ideas and priorities for the DietWise digital tools, specifically for Recipe Watch (RW). Rather than validating pre-defined solutions, the sessions focused on jointly exploring needs, challenges, and expectations that could inform the development of user-relevant and practical functionalities. Participants were recruited from the adult population, with particular emphasis on individuals who regularly cook and make use of digital recipe sources. This approach ensured that discussions remained grounded in everyday food-related decisions and reflected realistic expectations regarding the role and usability of digital applications in daily routines.

A consistent workshop structure was proposed across all pilots and adapted locally as needed, allowing outputs to remain comparable while accommodating differences in local context and group dynamics. Each session was structured as follows:

1. Introduction and context setting
2. Validation and extension of needs and barriers
3. Co creation of solutions through guided brainstorming and prioritization
4. Wrap up, next steps, and immediate feedback collection

To anchor discussions in concrete evidence rather than speculative debate, participants were invited to react to already gathered insights and data. They were encouraged to confirm what resonated with their own experiences, challenge assumptions that did not align with their reality, and highlight missing elements. This process helped refine and sharpen the understanding of key needs and barriers before moving into solution-oriented discussions. The co-creation phase focused on translating identified needs and constraints into feasible app functionalities, content formats, and communication approaches for RW. Proposed ideas were prioritized using light-touch methods such as voting or ranking, allowing participants to indicate features and approaches they considered most valuable. Throughout the workshops, local partners moderated the discussions and documented outputs using a structured template. This ensured consistent capture of participant profiles, expressed needs and barriers, proposed solution ideas, and their implications for the ongoing development of the DietWise tools.

2.1.3. Vulnerable Citizens

The vulnerable citizen workshops were designed as co-creation discussions. Participants were invited to help shape ideas and priorities for DietWise tools, including RecipeWatch. The sessions aimed to ensure that the perspectives of groups facing specific social, economic, or age-related challenges were meaningfully integrated into the development process by recruiting citizens from disadvantaged backgrounds, thereby supporting more inclusive and accessible ICT solutions. Participants were recruited from the adult population with specific elements – as the definition described in the context of D2.4 – of vulnerability based on the target groups of the pilot programs that will be implemented within the next months in the DietWise project. Hence, parents and school staff from areas of low socioeconomic status, elderly persons as well as minorities were the main persons who participated in the co-development sessions.

A consistent four-phase structure – as in the co-development sessions with citizens – was proposed and adapted locally, keeping outputs comparable while allowing partners to reflect local context and group dynamics. To support grounded and meaningful discussion, participants were asked to react to already collected evidence and data, confirm what resonated, challenge what did not, and add what was missing. This was used to sharpen the list of needs and barriers before discussing solutions. During the co-creation phase, participants worked together to translate identified needs and constraints into practical suggestions for app functionalities, content formats, and communication approaches for Recipe Watch. The proposed ideas were then prioritized using light-touch techniques, such as voting or ranking, allowing participants to highlight the features and types of support they perceived as most relevant and valuable in their daily lives. Across all workshops, local partners facilitated the discussions and systematically documented the outcomes using a structured template. This approach ensured consistent recording of participant profiles, articulated needs and barriers, proposed solution concepts, and their implications for the continued development of the DietWise digital tools.

2.2. PILOT SPECIFICS AND RESULTS OF CO-DEVELOPMENT ACTIVITIES

In the following subsections, we describe the pilot-specific adaptations and summarize the key insights, findings, and outcomes derived from the co-development workshops, structured by pilot country and stakeholder group as described in deliverable D4.1.

2.2.1. Influencers

Belgium

The Belgian co-creation workshops engaged a small but strategically relevant group of influencers (n = 3), including a dietitian, a recipe and lifestyle platform owner, and a recipe platform owner with ages varying from 30-ish to 55 years old. Discussion focused heavily on workflow integration and reputational risk, rather than on influencer-style content creation.

Participants highlighted following needs and requirements:

- Simple tools that integrate into existing workflows, including legacy systems and editorial processes.
- Scalable automated nutrition intelligence, including automatic nutrition analysis.
- Evidence of reliability, including validation data and practical proof that suggestions are workable and taste acceptable.
- Flexibility, with partners able to choose where and how to apply optimization rather than being pushed into an all or nothing approach.

Additionally, the discussion surfaced several barriers and risks:

- Previous exposure to AI generated recipes has produced unrealistic quantities, odd combinations, and substitutions that fail in practice, creating a credibility risk.
- Concerns about overcorrection, monotony, and messaging that feels moralizing, with strong sensitivity around the term “diet”.
- Tension between dietary guidelines as pattern level advice and the feasibility of enforcing those expectations at individual recipe level.
- Scale problems with human oversight, fear of losing control over brand, traffic, and mission alignment, particularly where local agriculture mandates apply.

In terms of opportunities and partnership models participants explored the following options:

- RCA positioned as an internal back-office service that strengthens partners’ expert image without heavy staffing demands.
- Process optimization through prepublication checks within recipe development workflows.
- Cautious interest in a quality stamp and related search filters, provided criteria and communication are transparent and applied at meaningful scale.
- Feasible partnerships focused on embedded services and long-term collaboration rather than one off campaigns.

Based on the above discussions, several concrete refinements to the RCA concept were proposed. These included the introduction of seasonality and locality filters, variation in suggested modifications, configurable levels of strictness, transparent validation indicators, and co-designed label logic. At the same time participants clearly articulated boundaries for acceptable development directions, emphasizing the importance of avoiding the creation of a competing consumer platform and cautioning against recipe level enforcement of pattern level guidelines.

Greece

The Greek workshops involved a mix of food influencers and bloggers (n=3), and dietitians (n=2). While several participants did not perceive the Responsible Cooking Alliance (RCA) as directly relevant to their own professional activities they expressed greater interest in opportunities linked to RecipeWatch, particularly in terms of visibility, collaboration, and promotion through project-related channels.

Across discussions, participants mentioned the following concerning needs, barriers and opportunities raised:

- Strong emphasis on recipes being easy to execute, good tasting, and time efficient as prerequisites for healthier and more sustainable choices.
- Acknowledgement that some popular content categories are not linked to nutrition and sustainability expectations, which shapes what they believe audiences will accept.
- Motivation to associate with a research initiative, connect with collaborators, and gain promotion opportunities through project channels and tools.

Within this context, usability and intuitive design were treated as non-negotiable. Participants also raised recurring concerns about ingredient substitutions that do not work in practice, drawing on their own experimentation experience and highlighting the importance of practical cooking feasibility alongside nutritional considerations. In terms of implications for the DietWise tools, influencers indicated that adoption would depend on a combination of factors, including the overall quality of the user experience, the credibility of taste and sensory outcomes suggested by the tools, and the presence of tangible value for them as content creators. Such value was frequently framed in terms of promotion and recognition.

Lithuania

The Lithuanian workshops engaged female influencers and nutrition-related content creators (n=2+1+1) aged between 28 and 36, combining non-nutrition-trained bloggers with a qualified nutritionist. This mix of profiles shaped discussions around perceived relevance, confidence, and appropriate positioning of nutrition-related tools. Participants expressed following needs and barriers:

- A strong preference to avoid one-sided content, maintaining space for both healthy and indulgent recipes, tied to personal values and audience expectations.
- Healthy content was described as less engaging than taste focused content, with concerns that strict health framing would feel preachy and alienate followers.
- Confidence barriers for non-qualified influencers who do not want to appear out of place when communicating scientific information.
- Mixed perceived value of RCA depending on education and role, including a view that a qualified nutritionist may not need it.
- Strong scepticism towards a points-based scoring system, seen as too extreme and not reflecting nuance, with a preference that any score should be visible only to the creator.

Despite the above reservations, participants recognized potential value in increasing the presence of healthier nutrition content and expressed interest in participation in an EU-funded project and related collaboration opportunities. However, uncertainty remained regarding audience interest and engagement with such content. Similar to the Greek pilot, participants stated they could not clearly understand RCA without a demo. They expected multiple useful functions, potentially commodifiable, such as nutrition calculation or shopping list tools, and advised against a didactic, overly scientific tone.

Cross-pilot takeaways that matter for development

Across pilots, trust and credibility emerged as critical gatekeepers for adoption, particularly in contexts where professional reputation, brand identity, and institutional mandates are at stake. The RCA was consistently perceived as more acceptable when positioned as an embedded back-office capability, rather than a standalone, public facing platform competing for traffic, ownership, or audience attention. Participants across countries identified overcorrection and moralizing framing as high-risk factors, both in terms of audience reception and partners' willingness to associate with the tool. Practical cooking feasibility – especially regarding ingredient substitutions and realistic portion sizes- was repeatedly highlighted as being as important as scientific or nutritional alignment. Influencers also clearly communicated expectations of tangible value for participation, often framed in terms of visibility, collaboration opportunities, or tools that reduce workload, rather than alignment messaging alone. An important limitation across several discussions was the absence of a working demo, which constrained participants' ability to assess usefulness and resulted in some feedback remaining hypothetical. This limitation was particularly evident in the Greek and Lithuanian pilots, where participants explicitly requested a demonstration, in order to properly evaluate the relevance and potential value of the RCA. A consolidated overview of the cross-pilot takeaways is provided in Table 1.

2.2.2. Citizens

Belgium

In Belgium, participants (n=6 adults across 2 sessions aged 35 to 66) stated that they mainly use a wide range of cooking applications and recipe websites, but generally reported limited engagement with influencers unless they were well-established culinary figures. While participants expressed openness to trying new digital tools, this willingness was strongly conditional on ease of use and the avoidance of a paternalistic or judgmental tone.

Key needs raised by Belgian citizens centered on:

- Speed and simplicity, especially for busy households.
- Ingredient substitution, leftover management, and seasonal and local filtering as practical routes to healthier and more sustainable cooking.

At the same time, participants identified several barriers, including:

- information overload caused by the proliferation of digital tools,
- a perception that healthier and more sustainable food options are more expensive,
- limited confidence in adapting recipes, and
- persistent time pressure.

When discussing potential solutions, participants prioritized:

- Personalization filters for allergies, dislikes, and dietary preferences.
- Meal preparation and freezing guidance, plus short visual demonstrations for techniques.
- Layered information, with simple scores or cues upfront and deeper explanations available when users want them.

In terms of implications for Recipe Watch, Belgian participants were explicit that a new application must be immediately intuitive and maintain a non-judgmental tone to sustain engagement. They also advised against explicit diet framing in user-facing language and suggested introducing sustainability concepts through local and seasonal options, which are perceived as more achievable and less prescriptive.

Greece

Greek participants (n = 6 adults aged between 30 and 65) reported frequent exposure to recipe content through social media platforms, particularly Instagram and TikTok. They described differing patterns of recipe discovery, with women more likely to try recipes encountered passively through social media feeds, and men more likely to search deliberately for specific recipes. Nutritional guidance was typically applied informally, rather than through structured use of dietary guidelines.

Following needs and barriers emerged:

- Cost and portion flexibility were central, including the ability to adapt quantities without complex measurements.
- Time constraints varied by life stage, but convenience still strongly influenced choice.
- Sustainable eating was often equated with plant-based cooking, with mixed acceptance and some concerns about satiety.

Proposed solution ideas and priorities focused on:

- Suggesting alternative recipes, not only alternative ingredients, especially options that are more time and cost efficient.
- Providing feasible meal preparation guidance and leftover tips linked directly to recipes.
- Offering cooking method alternatives and clear cautions where substitutions may affect outcomes.

The implications for RecipeWatch drawn from the Greek sessions point to convenience as the most credible pathway to supporting improved food choice. Participants also highlighted household realities, including the value of child-friendly alternatives to avoid the need to prepare separate meals for different family members.

Lithuania

Lithuanian participants (n = 6 women aged between 29 and 60) described digital recipe sources primarily as a source of background inspiration rather than something they follow step by step. They placed stronger emphasis on ingredient quality and provenance, with many expressing a preference for local foods and farmers' markets, and some noting that they grow their own produce.

Key needs and barriers included:

- Information overload and contradictions about what is healthy.
- Distrust in product quality, focused on how food is produced and what is added.
- Sustainability was often associated with organic products, while some participants also expressed strong distrust of organic claims.
- Healthy eating guidelines were not consciously used and were often interpreted through personal health experiences and restrictions.

Following solutions and priorities were discussed:

- Nutritional value calculation as a core function.
- Ability to enter available ingredients and receive recipe suggestions aligned with what is already at home.
- Notifications to support continued use.
- Suggestions linked to individual nutritional needs, including simple explanatory guidance where relevant.

In terms of implications for RecipeWatch, Lithuanian participants questioned whether actively checking the healthiness of recipes would feel natural or would be perceived as an additional step in an already demanding cooking process. They suggested that a single-function tool would not be sufficient, and that value would come from combining recipe support with nutrition calculation and personalized decision support.

Cross pilot takeaways that matter for development

Across all pilot countries, ease of use consistently emerged as a precondition for adoption rather than a user preference (see Table 1). Participants repeatedly emphasized the importance of guidance that feels practical and respectful, rather than corrective or moralizing. Layered information was identified as a promising compromise, offering quick cues for most users while allowing access to deeper detail for those who seek it. Finally, household realities were central across contexts, including the need for portion scaling, leftover management, and child-friendly alternatives.

2.2.3. Vulnerable Citizens

Belgium

The Belgian workshops involved adult women (n=6) of mixed ages, ranging from approximately 30 to 60 years old, primarily with a Turkish cultural background (n=5) and one participant with a Belgian cultural background. Some of the participants had professional experience related to food and health. All participants reported regular use of smartphones and multiple applications and were familiar with searching for recipes online and via social media. Most participants followed religious food rules, including halal requirements and restrictions such as no pork or alcohol.

Discussions highlighted following needs and barriers:

- Accessible, culturally familiar guidance: Strong preference for short (2–3 min) recipe videos featuring recognizable Turkish and Flemish comfort foods—familiar, filling dishes over "fancy" recipes—to balance emotional comfort with healthier eating.
- Healthy swaps without complexity: Demand for tasty, lower-calorie alternatives (especially desserts) and simple fat-reducing techniques (baking vs. frying, yoghurt-based sauces), without requiring precise measurements or complex steps.

- Personalization within real constraints: Interest in tailoring to tastes, dislikes, and dietary restrictions, but with recipes that work "by feel" and don't rely on weighing or extensive input.
- Cost and availability as major barriers: Healthy/special products (e.g., gluten-free) are perceived as expensive and hard to find; cultural habits (bread with every meal, high fat use) and budget pressures make portion control and change difficult.

Solution ideas and priorities:

- There is a strong need for healthy, affordable and culturally adapted recipes, including attention to e.g. halal rules and gluten-free options.
- Participants prefer practical advice on portions, ingredients and cooking methods.
- Trust and understanding are central: users want to know why a suggestion is healthier, in simple language.

Implications for RecipeWatch include a strong emphasis on short, actionable suggestions, such as one-line substitutions while avoiding long blocks of text. Robust personalization and filtering options are essential, particularly for religious rules (halal, no alcohol, gelatine), gluten-free diets, and cultural cuisine preferences. A budget-friendly dimension was identified as critical, including highlighting cheaper alternatives and recipes using accessible ingredients. Participants also suggested the inclusion of an ingredient-based recipe finder ("What can I cook with what I have?"). Explanations supporting health-related swaps should remain simple but meaningful, reinforcing key concepts such as fibre intake, satiety, and fat reduction over time to support learning.

Greece

Greek participants (n= 6 split over two workshops of 4 and 2 individuals. Adults aged 35-55 years old, educators and low-income parents/caregivers) reported frequent exposure to recipe content via social media platforms, applications, and websites. Digital literacy was not identified as a barrier, as participants already regularly search for recipes online, primarily when looking for something new rather than for everyday cooking.

Key needs and barriers raised by participants are:

- Cost and time constraints: Participants consistently cite limited budgets and lack of time as the main barriers to home cooking.
- Ingredient practicality: Recipes using expensive, non-reusable, or perishable ingredients (e.g. quinoa, flax seeds, avocado) discourage experimentation, especially when leftovers are hard to store or reuse.
- Recipe complexity: Dishes that require many steps, tools, or ingredients are seen as impractical and reduce motivation to try healthier options.
- Not a digital access issue: Digital literacy and access to apps or online recipes are not barriers; participants already search for recipes online, mainly when trying something new.

In terms of solution ideas and priorities following points were raised:

- Practical ingredient substitutions: Provide easy, familiar alternatives for sugar, flour, butter, or missing ingredients, prioritizing options users already have at home or can easily find, to increase acceptability and reduce cost and effort.
- Flexible cooking-method guidance: Offer clear method swaps with time/temperature equivalences (e.g. oven vs. air fryer, boiling vs. steaming) so users can adapt recipes to their equipment and time constraints.
- Meaningful personalization: Allow personalization by allergies, taste preferences, and dietary needs, plus dynamic priorities (e.g. saving time on weekdays, saving money at month-end) using a flexible, context-dependent checklist rather than fixed settings.
- Transparency, trust, and continuity: Position the app as a supportive companion by explaining the reasoning behind recommendations when needed, optionally sharing its evidence base, and enabling users to save, like, and revisit recipes and adopted suggestions in a personal library.

Implications for RecipeWatch highlight that day-to-day concerns seem to be more critical to perceived usefulness of the application. Additionally, recommendations should not compromise recipe success or taste. On the other side, recommendations extending beyond environmental or plant-based criteria, should ensure that alternatives are anchored to restrictions that don't compromise on the recipe sensory outcome.

Lithuania

In Lithuania, participants (n = 9) were elderly individuals – primarily women, all with higher education backgrounds across different fields. Participants were digitally literate, socially active and familiar with online recipes, though they tend to use them mainly for social occasions rather than daily cooking.

The needs and barriers identified are the following:

- Health-first decision making: Perceived healthiness and ingredient composition are the most important factors in food choices, outweighing concerns about price or taste.
- Structural distrust as a major barrier: The food industry is viewed as the main threat to healthy eating, with strong beliefs that foods are contaminated with chemicals and that individual choices have limited impact compared to industry influence.
- Knowledge and affordability gaps: Lack of nutritional knowledge and high prices are seen as practical barriers to eating healthily and sustainably.

Solution ideas and priorities discussed:

- Positive reception to real-time guidance: Participants responded well to the idea of an app offering healthier alternatives, especially when suggestions are short, concrete, and easy to apply (e.g., reducing salt/sugar/fat or switching from frying to oven cooking).
- Health-driven substitution acceptance: Willingness to accept both familiar and less familiar substitutions is primarily motivated by clear health benefits; however, resistance remains for changes with noticeable taste or texture impacts (e.g., replacing white flour with whole-grain flour).
- Selective need for support: Some substitutions (e.g., whole-grain pasta instead of regular pasta) are seen as obvious and do not require app-based guidance, while others benefit from targeted recommendations.
- High trust with transparent explanations: Participants are comfortable sharing data for personalization and generally trust app recommendations, provided they are supported by brief, scientifically grounded explanations delivered in an expert tone rather than vague claims.

For RecipeWatch, implications include placing greater emphasis on health-related aspects than on sustainability, as the concept of sustainability appears to be less clear to older adults. Additionally, the app should include elements that help the elderly gain more knowledge about healthy eating and national nutrition guidelines as well as cooking skills. Lastly, the app should highlight how changes affect health, as this appears to be the most important factor for the elderly. However, it should be noted that the sample group consisted of educated seniors with likely higher-than-average income levels.

Cross pilot takeaways that matter for development

Across all pilot countries, cost, time, ingredient availability, and recipe simplicity strongly influence adoption (see Table 1). RW-features should prioritize affordable and reusable ingredients, minimal preparation steps, flexible cooking methods, and “cook with what you have” functionality, while maintaining taste and recipe success. Participants consistently favoured concise, and short-format content, including brief videos, over long texts or complex instructions. Personalization emerged as a consistent requirement, covering dietary restrictions, allergies, cultural or religious rules, taste preferences, and dynamic priorities related to time, cost, and health. Privacy concerns were generally low; however, personalization must remain lightweight and easy to adjust. A high baseline level of trust exists across contexts, provided the app maintains an expert, evidence-based tone and positions itself as a supportive companion, allowing users to accept, ignore, or save suggestions. Across all pilots, health was the primary driver of acceptance, with nutritional value, ingredient compositions, and clearly articulated health benefits consistently needing to be foregrounded.

Stakeholder group	Tool	Key proposals & value drivers	Core design & functional requirements	Risks, constraints & watchpoints
Influencers	**RCA**	• Position RCA as a **supportive, back-office capability**, embedded in existing workflows rather than a public-facing platform • Offer **tangible participation value** (visibility, collaboration, workload reduction) •Emphasize **practical cooking support**, not only nutritional optimization	• **Seamless integration** with existing content creation processes • **Realistic** ingredient substitutions and portion guidance • **Non-moralizing, non-corrective** framing of recommendations • **Clear articulation of benefits** for professional use	• Risk of rejection if perceived as competing with influencer brands or audiences • Overcorrection or moralizing tone may harm credibility and partnerships • Lack of a working demo limits ability to assess usefulness and relevance
Citizens	**RW**	• Enable **informed, low-effort decision-making** in everyday cooking • Support **household realities** and flexibility rather than idealized diets	• High **ease of use** as a baseline requirement, not an enhancement • Respectful, supportive **guidance** tone • **Layered information** (quick cues + optional detail) • Portion scaling, leftover management, and child-friendly options	• Perceived judgement or correction may reduce trust and engagement • Overly dense information risks disengagement
Vulnerable citizens	**RW**	• Prioritize **health, affordability, and feasibility** as primary adoption drivers • Position RW as a **supportive companion**, not a directive authority	• **Affordable**, widely available, and reusable ingredients • Minimal preparation steps and **flexible** cooking methods • "Cook with what you have" functionality • **Short-format** content (e.g. brief videos) • Lightweight, **adjustable personalization** (dietary, cultural, time, cost, health needs)	• Complexity or high cognitive load may exclude users • Personalization that feels rigid or intrusive could undermine trust • Health benefits must be clearly articulated to sustain engagement

Table 1: Cross-pilot takeaways that matter for development (Belgium, Greece, Lithuania)

3. EXPERT-DRIVEN CO-DEVELOPMENT AND ADAPTATION

3.1. ALIGNMENT WITH NATIONAL GUIDELINES

Matching the ICT solutions with nutrition guidelines was approached first and foremost as a nutritional development and harmonization exercise: before deciding how digital tools should filter, score, or recommend recipes, we needed to define what "healthy and guideline-compliant" *means* in a way that is evidence-based, transparent, and adaptable across countries. VU, VIGL, IHU, PROL, and PHB jointly contributed to the co-development of a rule set that can be implemented in ICT solutions while remaining aligned with national nutrition policies.

The first pillar of the development was establishing a shared scientific rationale for prioritizing dietary targets. For this, we used evidence from the Global Burden of Disease (GBD)[1] study to identify the dietary risk factors most consistently associated with morbidity and mortality (e.g., low intakes of protective food groups such as fruits, vegetables, whole grains, legumes, nuts/seeds; and high intakes of risk-associated components such as excess sodium and certain ultra-processed or high–saturated fat/sugar products, depending on country framing). This step ensured that our criteria were not only "policy compliant" but also anchored in widely recognized population health evidence, helping us justify why specific recipe rules and thresholds were selected.

A second pillar was a structured inventory and comparison of Food-Based Dietary Guidelines (FBDGs) across all participating countries, as documented in D2.1. This process focused on identifying common parameters within National Food-Based Dietary Guidelines (FBDGs) and, where available, food-related sustainability recommendations. Nutrition experts reviewed nutritional and sustainability guidelines across all participating countries, including accompanying visual models where applicable, such as the Food Triangle[2] in Flanders. These FBDGs were then compared to the map:

- Common cross-country principles (e.g., more plant-based foods, variety, water as default beverage, limiting salt/sugary foods, preference for minimally processed foods).
- Areas where guidance differs in emphasis or terminology (e.g., specific recommendations on dairy, meat frequency, oils/fats, wholegrain definitions, discretionary foods).

This comparison step was essential because ICT solutions typically operate on structured inputs (ingredients, quantities, portions, nutrients), while national FBDGs are often written for public communication. Our goal was therefore to translate FBDGs into implementable, auditable criteria that can be encoded as recipe rules.

Within this approach, the Zeker Gezond selection criteria and its grounding in the Flemish Food Triangle were treated as a useful reference baseline, particularly for demonstrating how guideline principles can be operationalized into concrete, implementable criteria for recipe selection. Based on the overview and level of detail obtained from this exercise, several concerns were identified regarding the granularity, specificity, and recipe-level applicability of FBDGs. By design, FBDGs are generally intended to guide overall dietary patterns rather than to inform recipe-specific evaluations or modifications. In addition, the selection and harmonization of guidelines posed challenges in a limited number of cases due to inconsistencies across countries.

Based on these considerations and following the agreement that the nutritional rules underpinning public-facing ICT tools should be evidence-based, a decision was made to structure the rules-based system on the GBD Study. This decision was based on three main factors:

(i) The GBD Study provides quantitative estimates of the health impacts associated with a defined set of specific dietary factors,
(ii) these dietary factors can be readily leveraged for recipe-level association and to generate recommendations; and
(iii) it serves as the evidence base for the Belgian national FBDGs, thereby offering a robust and well-established foundation to be used as a proof of concept for the technical development of the ICT tools.

3.2. PERSONALIZATION

The personalization engine of RW utilizes a data-driven approach to align individual dietary habits with evidence-based nutritional standards. This mechanism is primarily grounded in the GBD study, incorporating dietary risk exposure estimates to tailor interventions. The GBD Study explicitly accounts for age and gender differences by recognizing that individuals at different stages of life, and of different sexes, have distinct nutritional needs and health risks. To reflect

[1] Institute for Health Metrics and Evaluation (IHME). Global Burden of Disease 2023: Findings from the GBD 2023 Study. Seattle, WA: IHME, 2025.
[2] Flemish Institute for Healthy Living 2021. Eating according to the food triangle: good for yourself and the planet. In cooperation with the Department of the Environment and the Flemish Agency for Care and Health. Laken (Brussels)

this, the study estimates the health effects of specific dietary factors separately by age group and gender. In addition, the GBD Study defines age- and gender-specific levels, representing the intake levels of dietary factors associated with the lowest health risk for each population group. Together, these elements allow dietary risks and benefits to be assessed more accurately across population groups, rather than applying a single "one-size-fits-all" optimal intake level. By structuring the nutritional rules on this same framework, the recipe recommendations could be personalized according to the user's age and gender, ensuring that suggested recipe adaptations and prioritization are aligned with population-specific health evidence.

In that sense, personalization in the ICT solutions is achieved by segmenting users into specific demographic profiles based on gender and three distinct age groups (see Annex Figures Figure A 3Figure A 4Figure A 5). The system applies a rules-based logic to provide recipe recommendations that are prioritized according to these parameters. By mapping individual profiles against a list of dietary factors that impact human health (e.g. the intake of fruits, legumes, and processed meats), the application executes food-level substitutions. These substitutions are triggered by specific ingredients within a recipe and replaced with alternatives that improve the recipe's alignment with the country-spread guidelines that were defined.

While co-creation activities highlighted strong expectations for more granular personalization (such as accounting for allergies, intolerances, cultural or religious dietary rules, taste preferences, time constraints, or budget considerations) these dimensions were not included in the current version of the personalization engine. This decision reflects a deliberate design choice to prioritize robustness, transparency, and public health relevance in an early-stage implementation. Age and gender were selected because they are consistently available, require minimal user input, pose low privacy risks, and are directly supported by strong epidemiological evidence within the GBD framework. In contrast, more individualized dimensions often require extensive self-reporting, complex rule sets, or additional validation mechanisms, which could increase system complexity, user burden, and data protection requirements.

Despite these limitations, the current form of personalization does also address several needs of vulnerable citizens in an indirect but meaningful way. By grounding recommendations in age- and gender-specific health risks and emphasizing nutritionally adequate, evidence-based substitutions, the system supports population groups with higher baseline dietary risk. Moreover, the rules-based and transparent nature of the approach ensures predictable and explainable recommendations, which is particularly important for users with lower digital or health literacy. The emphasis on clear guidance, minimal required input, and avoidance of moralizing feedback further aligns with the needs expressed by vulnerable groups during the co-creation process. Future development of RW could progressively incorporate additional personalization dimensions, building on user feedback and technical validation.

3.3. NUTRITIONAL SCORING MECHANISM

To support informed decision-making and encourage healthier dietary choices, the RecipeWatch app introduces a nutritional scoring system that translates complex nutritional information into accessible and actionable format for its users. The score is presented directly within the app alongside each recipe, allowing users to quickly understand how a given recipe aligns with evidence-based healthy eating recommendations and to compare alternatives. Rather than functioning as a judgement or label, the score is designed as a decision-support feature that helps users reflect on nutritional quality and explore opportunities for healthier adaptations, providing actionable insights and fostering behaviour change. The recipe scoring system builds on the GBD dietary risk factors as the core set of components used to assess alignment with healthy eating. Each component corresponds to a dietary factor that is either encouraged or should be limited. By evaluating the presence or absence of these factors within a recipe, the system computes a composite score that reflects the recipe's overall nutritional profile. Higher scores indicate a greater contribution to a healthy diet and a lower associated dietary risk.

Regarding computation, each recipe is assessed against 15 GBD dietary components (see Annex Figure A 1). For encouraged components, a recipe earns points when the relevant healthy food group or nutrient is present. For limited components, points are awarded when the ingredient is absent or remains below a predefined threshold (e.g., low salt content or absence of processed meat). If a recipe lacks encouraged components, it does not receive the corresponding points. The individual component scores are summed to produce a raw score with a maximum of 15 points, which can be retained in its original form or scaled to any other. In effect, the score reflects the extent to which a recipe contributes to mitigating specific dietary risks, such as a "diet low in X", by providing or limiting the relevant components.

Partners acknowledge that any recipe scoring or rating system presents inherent methodological challenges and may raise valid concerns, particularly with respect to the implicit characterization or "labelling" of individual recipes. To address this, **a decision was made to provide two versions of the score: one representing the original recipe, and a second reflecting the potential outcome if the user adopts the AI-generated recommendations**. This approach shifts the emphasis from judging recipes to supporting informed choices and encouraging the uptake of healthier adaptations.

4. DESCRIPTION OF THE APPLICATIONS

4.1. RECIPE WATCH

RecipeWatch (RW) is a standalone mobile and web application addressed to citizens, aiming to support them in their nutritional decisions. Given a recipe, the application provides suggestions for ingredient alternatives towards a healthier and more sustainable recipe version. Citizens actively participate in the adaptation process, regarding suggestions by reviewing the proposed substitutions, accepting or rejecting them. Registered users enjoy a more personalized experience, where suggestions are shaped according to their age and gender. In order to promote transparency and user engagement, RW allows users to view aggregated information regarding which suggestions have been accepted or declined by other users for the same recipe. This feature was deliberately included as part of a behavioral intervention strategy, testing whether social and collective feedback can act as a nudge to support reflection and uptake of healthier choices, without prescribing or enforcing specific actions (see Deliverables D3.1-4 on behavioral interventions). This collective feedback contributes to the formation of a statistical knowledge base, which may -in future development phases- be exploited to improve the quality of recommendations and enable more refined user profiling, ultimately enabling more relevant, practical, and trustworthy suggestions for users, thereby improving their experience and confidence in the recommended adaptations.

Download & Sign up

A citizen who learns about the RW app will be able to download it on both Android and iOS devices. The application supports both registered and anonymous usage, allowing users to browse recipes without creating an account. In order to create an account, the user needs to input minimal personal data, thus ensuring privacy preservation and increasing user trust. Authentication is handled through email and password credentials. Additional information includes their name optionally, as well as the two mandatory attributes: date of birth and gender. The latter two constitute the personalization part and lead to user-specific suggestions as explained in the section 3.2. Registered users may also manage specific account settings. First of all, for security reasons they are able to update their password whenever they feel it is necessary. They may update their name, date of birth and gender, in case of wrong initial input or change of data where it is applied. Lastly, the users can select their preferred language of operation; at the current stage the application is available in English, with additional language options for the pilot countries foreseen in future development phases.

4.1.1. User Experience

RW is designed to be easy to use and uncluttered, providing information in a clear and accessible manner aligned with the end users' needs, particularly those of vulnerable citizens as expressed during the co-creation sessions. Given a recipe, RW analyzes its content, identifies health and nutritional gaps and subsequently gives suggestions for alternative ingredients. The user can then review each suggestion and decide whether to adopt or decline it. In that sense, the user gives explicit feedback to the application, which may be stored accordingly, in order to enable future scenarios of further shaping the food environment to user needs and/or refine the suggestions to provide more relevant advice for each user. Keeping such data can possibly lead to both better suggestions and better profiling of the users, thus more granularity in personalization. The user experience can be summarized into the following workflow:

The user identifies a recipe online, copies and pastes the link into the app UI. The app extracts the recipe in the background and RW displays only the clean, extracted recipe including ingredients and instructions. After the recipe is properly extracted, an LLM is invoked to provide suggestions. The latter are then presented to the user overlaid on the extracted recipe. For each proposed substitution, the user is also provided with aggregated statistics indicating how many other users have accepted/rejected the same suggestion. This way the user benefits from the collective experience and builds trust through transparency. Once the user finalizes their choice, the changed recipe is clearly indicated through the UI, as shown in Figure 1, ensuring full awareness of the applied changes. As explained in section 3.2, suggestions are strictly limited to ingredient substitutions, resulting to healthier and/or more sustainable alternatives.

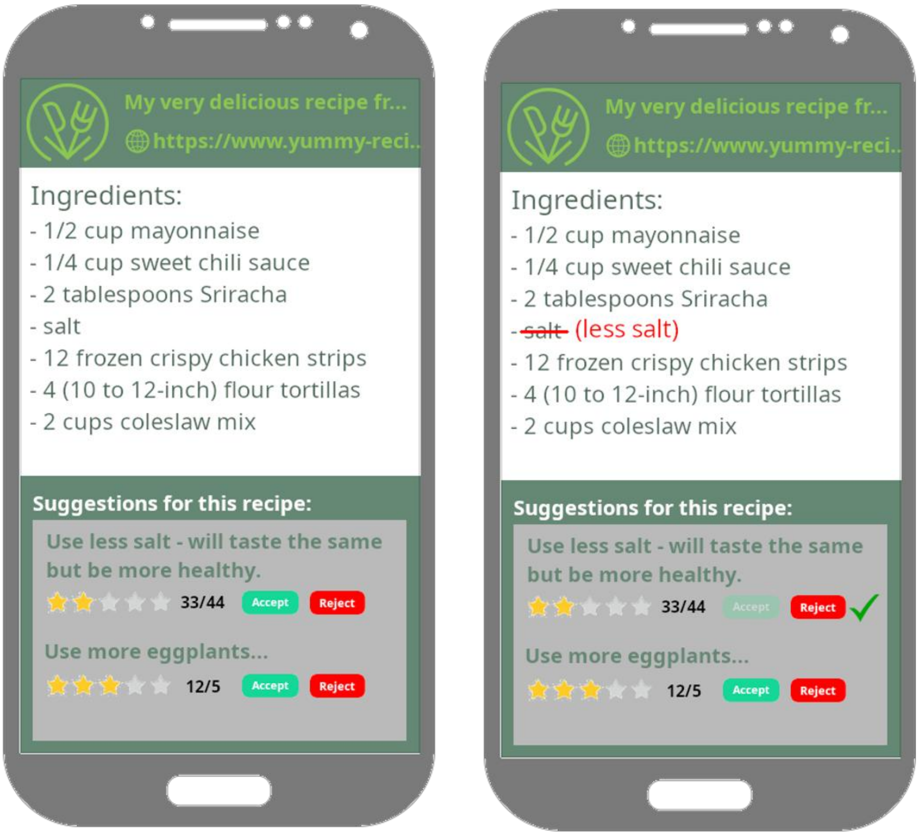

Figure 1: RW mock-ups with extracted recipe (left) and recipe with adoptions overlaid (right).

4.2. RESPONSIBLE COOKING ALLIANCE TOOL

Responsible Cooking Alliance (RCA) is a browser extension/plugin that interacts with cooking-related websites of influencers. Its primary objective is to support them in assessing the compliance of their content with nutritional guidelines. Influencers can use RCA to get advice on how to enhance their content. RCA is mainly delivered through a browser extension. The plugin reads the current recipe page and interacts with the infrastructure of DietWise (see section 4.3) to provide recommendations. A dedicated web page is also provided for account registration and management.

4.2.1. Sign up

An influencer who finds out about RCA may use the dedicated website to register. To create an account the user needs to input minimal information, namely their email and preferred password. As users, the influencers can manage their password whenever necessary for security reasons. Having an account, the user may configure its basic settings, which is to change the given information, update password and set the preferred language; at the current stage the browser extension is available in English, with additional language options for the pilot countries foreseen in future development phases.

4.2.2. User Experience

Registered influencers browsing their own website can access RCA directly through a dedicated browser extension button located in the top-right area of the browser interface, as shown in Figure 2. By activating the extension, a side-panel popup is displayed, containing all feedback and recommendations generated by RCA (see Figure 3). The plugin extracts the recipe (both ingredients and instructions), analyses it and finally displays the recipe together with the suggestions. These include general advice, as well as specific ingredient substitutions. The suggestions aim to assist influencers in improving the nutritional alignment and sustainability profile of their published recipes, without altering the original content unless explicitly chosen by the user.

The RCA will not display an explicit score or ranking for individual recipes. This decision is motivated by both strategic (derived from co-creation sessions) and technical considerations. From a strategic perspective, the objective is to avoid positioning RCA as a source of negative publicity for content creators, which could hinder adoption and engagement. In practice, this could also discourage participation, as poorly ranked content would be unlikely to be shared or promoted. From a technical perspective, recipe content published online is subject to frequent updates and modifications. Accurately tracking such changes and continuously recalculating rankings would introduce significant technical complexity. By refraining from publishing rankings, RCA retains greater flexibility to expand its functionalities and enables influencers to use the tool supportively not only for content they own, but also when browsing and assessing recipes on third-party websites.

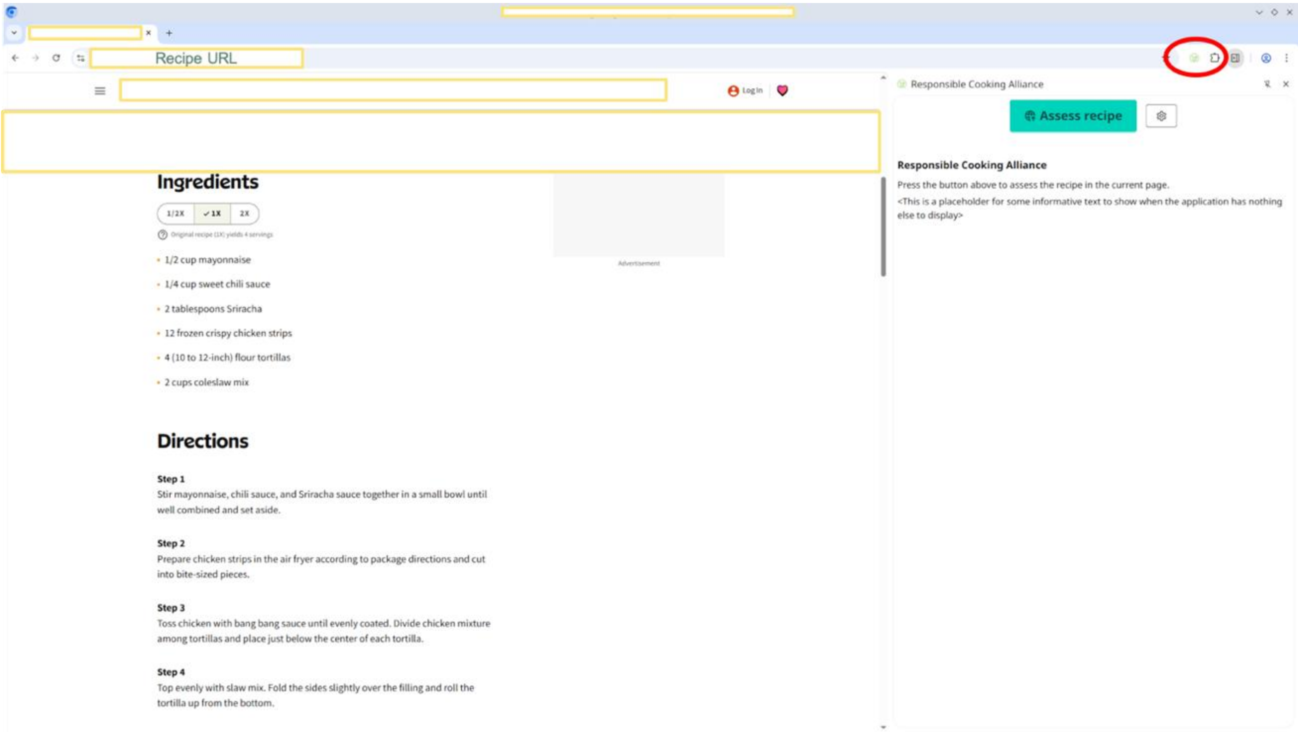

Figure 2: RCA browser extension; button activation.

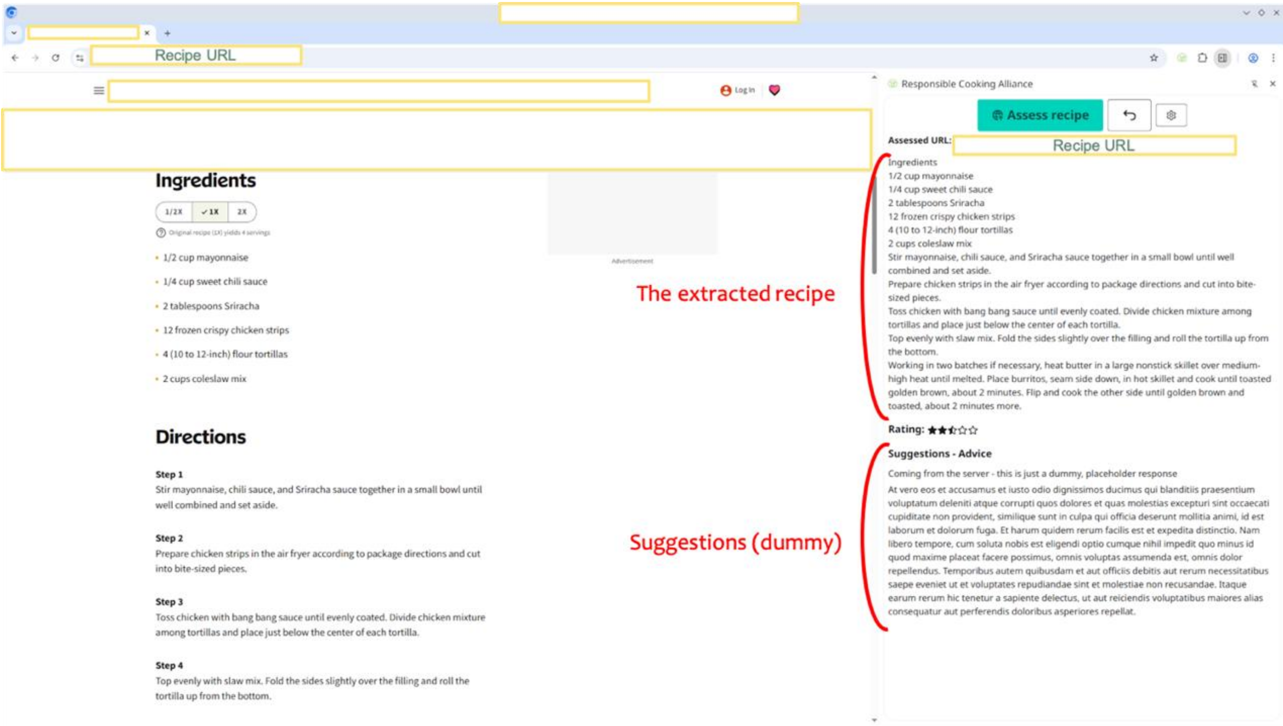

Figure 3: RCA side panel pop-up with dummy recommendations.

4.3. THE ARCHITECTURE OF THE ICT SOLUTIONS

This section describes the architecture of the ICT solutions. We first present a high-level overview of the system and gradually zoom in to describe the details of the containers. The two user-facing applications produced by DietWise, namely RW and RCA, share a common infrastructure that is implemented as services provided online to both of them (see Figure 4). Citizens can sign up with RW, while influencers with RCA. The content of RCA is, however, also visible to citizens, in the sense that they have access to recipes assessed by DietWise, which comply with nutrition guidelines. The account of every registered user -citizen or influencer- is stored in an Identity Management system.

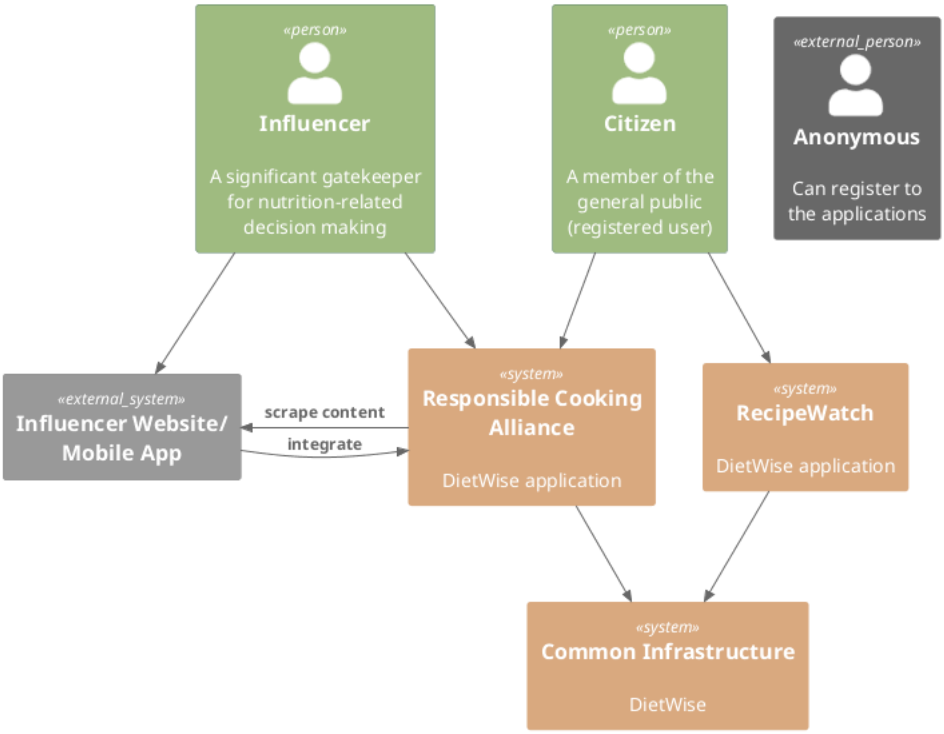

Figure 4: High-level overview of the system.

4.3.1. Technical Specifications

After presenting the high-level overview of the system and before proceeding to a more detailed architectural description, it is useful to briefly outline the tools and technologies adopted for its implementation. The implementation has been approached with a clear architectural vision, defining both what needs to be built and how the different components are expected to interact. Based on this understanding, development has already begun with the design and deployment of the core infrastructure supporting the system. The selected tools and technologies are well-known and industry-standard solutions, chosen to ensure reliability, scalability and ease of maintenance. It is important to note that the technical implementation is not considered static and may change according to the needs that arise as the system matures.

The chosen technologies are:

- Java 25 for the majority of the server-side components,
- Quarkus as the main server infrastructure,
  - NodeJS version 24 and Express services where needed (the latest long-term support version as of this writing)
- Ollama as the chosen infrastructure for running AI components,
- PostgreSQL version 18 for its general relational storage,
- Keycloak is the Identity Manager,
- Server components to be hosted in ICCS,
- Ionic 8 to develop cross-platform mobile applications,
- React 19 for all the UIs (both Ionic and the browser extension),
- TypeScript 5 for all NodeJS/browser/Ionic applications.

We strive to use the latest long-term support version for all components, where applicable. We upgrade regularly to stay up-to-date and mitigate any vulnerabilities or bugs in the infrastructure.

We incorporate the following tools in the development process:

- OWASP Dependency-Check through the Maven plugin,
- SonarQube for static code analysis.

4.3.2. DietWise Interoperability Framework and the Corresponding Architecture

For brevity the anonymous user (see Figure 4) is omitted and the description focuses on registered users, namely influencers and citizens**.** Figure 5 presents the main containers of the system under development and the interactions between them.

User registration, authentication and authorization is handled through the User App/Identity Management system (IDM). This component enables users to sign up with RW, authenticate in both applications, manage their credentials, and obtain appropriate access rights. It is implemented using Keycloak, which provides advanced identity and user administration capabilities. Authorized DietWise personnel may also use this system to register influencers (if self-registration is not enabled), register citizens when required, and deactivate or delete user accounts.

User related authentication data is stored in the User DB, a PostgreSQL database managed internally by Keycloak. This database contains personally identifiable information (PII) strictly related to authentication and authorization. It is intentionally kept separate from the application database to enhance data isolation and reduce privacy risks.

Responsible Cooking Alliance (RCA) is the browser extension/plugin used by influencers. It supports both Chrome and Firefox and is distributed through the respective browser extension stores. The extension authenticates users using the industry-standard OAuth2 authorization code flow with PKCE. As a browser extension, it has access to the content of the page currently visited by the user. This allows to send this content directly to the backend for assessing the recipe, after simplifying it as much as possible.

RecipeWatch (RW) is the mobile app intended for citizens. Developed using Ionic, it will be available on both Android and iOS platforms through the respective app stores. Like RCA, RW authenticates the user via OAuth2 code flow with PKCE, using the device's secure system browser tab. Unlike browser extensions, a mobile application cannot access the content of a mobile browser, even if the browser is embedded within the application. Therefore, the strategy for extracting the recipe from a web page is different than that of RCA. Instead of sending page content, RW sends the URL (address) of the recipe page to the server which is responsible for downloading the page, extracting/simplifying its content and assessing the recipe.

The DietWise Backend implements the core application logic of DietWise shared by both front-ends. This is the only server-side component directly exposed to the outside world. It is implemented using Quarkus and Java and acts as the orchestration layer of the system. It uses the Rendering Service and the AI to produce its results and send them to both front-ends. To assess a recipe, the backend coordinates the following workflow:

1. Download the recipe page content (the RCA performs this step directly; RW via the Rendering Service)
2. Minimize the content by erasing unrelated data/noisy content
3. Extract the recipe using either JSON-LD standardized format or AI (see Figure 7)
4. Feed the recipe, the dietary rules, user personalized information (applied only to RW) and other user preferences/past usage information to the AI
5. The AI produces the suggestions

Due to the computational cost of AI processing, recipe assessment may require several seconds. To improve user experience, the backend tries to progressively return partial results to the user as they are produced, thereby reducing perceived response time.

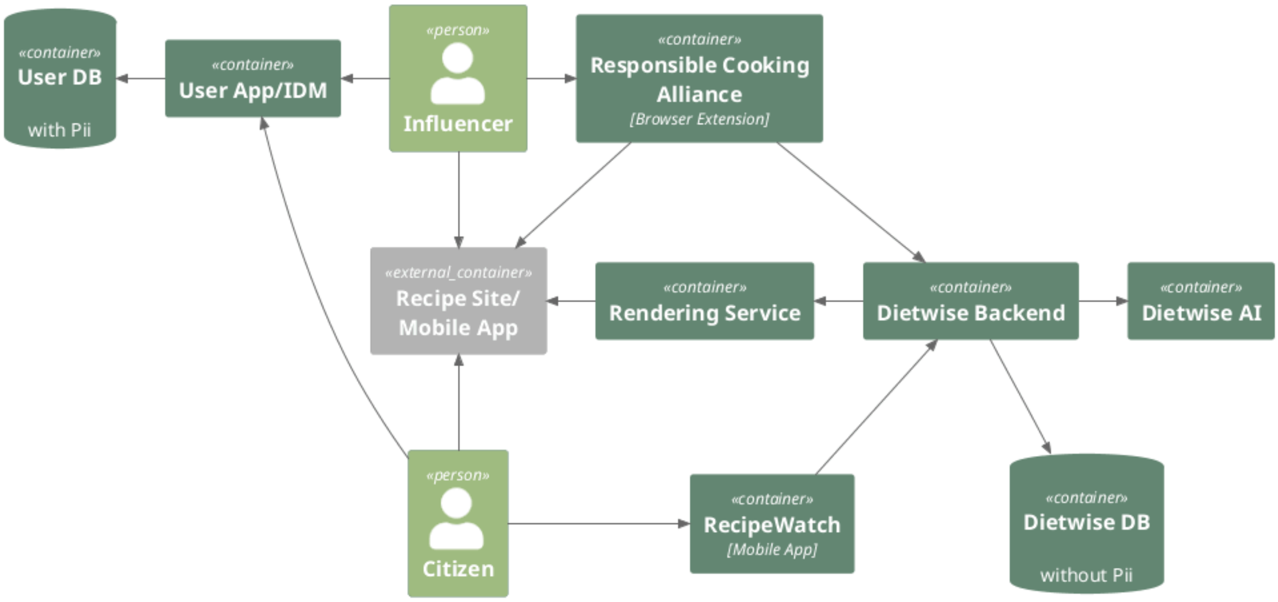

Figure 5: The main containers of the system and the interactions between them.

The Dietwise DB is the relational database (PostgreSQL) that stores all the information needed by the backend. Unlike the User DB, it contains no personally identifiable information. This ensures that stored data cannot be linked to specific individuals, even in the event of unauthorized access.

Dietwise AI constitutes the core analytical engine of the system. It implements the business logic, i.e. the rules provided by the dietician partners to assess a recipe and suggest alternative ingredients. The selected Large Language Model (LLM) is executed through Ollama. This component is internal and is not exposed to the outside world. The front-ends use it indirectly through the backend and no authentication is required within the trusted internal network.

Since RW cannot directly access the content of a web page, an additional service -the Rendering Service- has been implemented to retrieve and preprocess recipe pages on its behalf. The Rendering Service is implemented using Playwright, a library originally intended for testing web applications. Like Selenium it launches a headless browser (i.e. a browser that does not output to a screen, suitable for servers) and allows programmatic interaction with a page. We chose the NodeJS flavor of it and created the Rendering Service with Express 5. Like the AI, this service is not directly exposed to the external world. It is only accessible from the internal network. As such, it requires no authentication.

4.3.3. Recipe Extraction Approach

The first and most important aspect affecting both ICT solutions concerns the extraction of recipes from third-party websites. This is a critical step, since the quality of the extracted recipe directly impacts all subsequent processing stages. A major challenge arises from the heterogeneity of website structures. Most recipe websites utilize the JSON-LD standardized W3C format used to structure data, making it easily readable form machines and engines, which provides a well-defined and semantically structured representation of recipe information according to definitions from schema.org. In such cases, the extraction process is straightforward. All the information can be reliably identified and retrieved in a machine-readable format (JSON) from the page content. However, we have identified pages that do not follow this format properly. Common issues include inconsistent data (e.g., different instructions in the human-visible part of the page when compared to the JSON-LD), missing required or recommended fields, improper nesting and structure or incomplete schema implementations. Other websites rely on alternative, older formats, while others do not adhere to any structured data standard at all. In these cases, recipe extraction must be performed in a largely schema-agnostic manner, significantly increasing technical complexity and uncertainty. Since the DietWise applications are designed to operate on any publicly accessible website (excluding password-protected content), a robust extraction strategy capable of handling highly diverse page structures is required.

To address this challenge, the project adopts a hybrid extraction approach (see Figure 7**Fout! Verwijzingsbron niet gevonden.**) that combines JSON-LD-based parsing with Large Language Model (LLM)-based extraction, depending on the characteristics of each website. When reliable JSON-LD data are available, they are used as primary extraction source. When structured data are missing or unreliable, the LLM is tasked with identifying and extracting recipe information directly from the raw web content. This approach avoids the need for an extensive set of deterministic extraction rules, which would require a large number of site-specific scenarios, continuous maintenance and would be difficult to scale.

Nevertheless, LLM-based extraction introduces its own limitations. The quality of the extraction outcome is highly dependent on the quality and size of the input provided to the model. During preliminary experimentation, both successful extractions and failure cases were observed. Failures are primarily associated with websites that produce excessively long and noisy inputs, where recipe content is embedded among advertisements, navigation elements, and unrelated text. As a mitigation strategy, the extraction pipeline (in cases where JSON-LD cannot deliver reliable results) incorporates a two-step preprocessing approach:

1. Aggressive cleaning and simplification of the HTML content,
2. Conversion of HTML to Markdown, and
3. Submission of the resulting reduced content to the LLM for recipe extraction.

This strategy aims to minimize input length and noise, thereby increasing the likelihood of correct recipe identification. The current prompt for recipe extraction was derived through an iterative experimentation process involving the design, testing, and refinement of multiple prompt variants, with the goal of maximizing extraction accuracy and consistency both content-wise and regarding the output format. The optimized prompt used for this extraction process is shown in Figure 6. After evaluating several openly accessible LLMs, we selected nemotron-mini (4B) for deployment. This model is accessible via Ollama and is optimized for role-based interactions, retrieval-augmented generation (RAG) question-answering, and function calling. It supports a context window of 4,096 tokens and is released under a license suitable for commercial use. Nemotron-mini currently offers an effective balance between performance and computational efficiency for the recipe extraction task. Larger models may be considered in future iterations, subject to infrastructure capacity and performance requirements.

```
You are a recipe extraction engine.
Input is Markdown converted from a web page.
Extract one recipe only (the primary/complete recipe on the page).

Return STRICT JSON ONLY, with this schema:
{
    "name": string | null,
    "recipeYield": string | null,
    "recipeIngredients": string[],
    "recipeInstructions": string[]
}

Rules:
- Output must be valid JSON with double quotes and no trailing commas.
- Do NOT wrap the JSON in markdown or backticks.
- Ingredients and steps must each be single-line strings.
- Normalize ingredient/step text into plain strings (remove bullets, numbering, extra whitespace).
- Ignore navigation, ads, stories, comments, and unrelated content.
- If no recipe is present, return: {"name": null, "recipeIngredients": [], "recipeInstructions": []}
```

Figure 6: Tailored LLM prompt for recipe extraction in cases where JSON-LD cannot deliver reliable results.

Despite the adoption of a hybrid extraction strategy and multiple mitigation measures, there remains the possibility that extracted recipes may be incomplete, partially incorrect, or misaligned with the original content, particularly for highly unstructured or atypical web pages. Such inaccuracies could negatively affect the quality of subsequent analyses and recommendations, as well as user trust in the system. To address this challenge, future versions of the application will provide an explicit user-in-the-loop fallback mechanism. Users will be able to manually copy and paste recipe content into the application and, when necessary, review, correct, or supplement automatically extracted recipes. This functionality allows users to intervene in cases of faulty extraction, ensuring that the recipe data used for processing accurately reflects the intended content, thus promoting transparency and building user trust.

Beyond risk mitigation, this feature also represents an added-value functionality. It enables users to input their own original or family recipes, which may not be available online, and receive personalized recommendations and analyses based on them. In this way, a potential technical limitation is transformed into a robustness-enhancing mechanism that improves system reliability, user engagement, and overall flexibility.

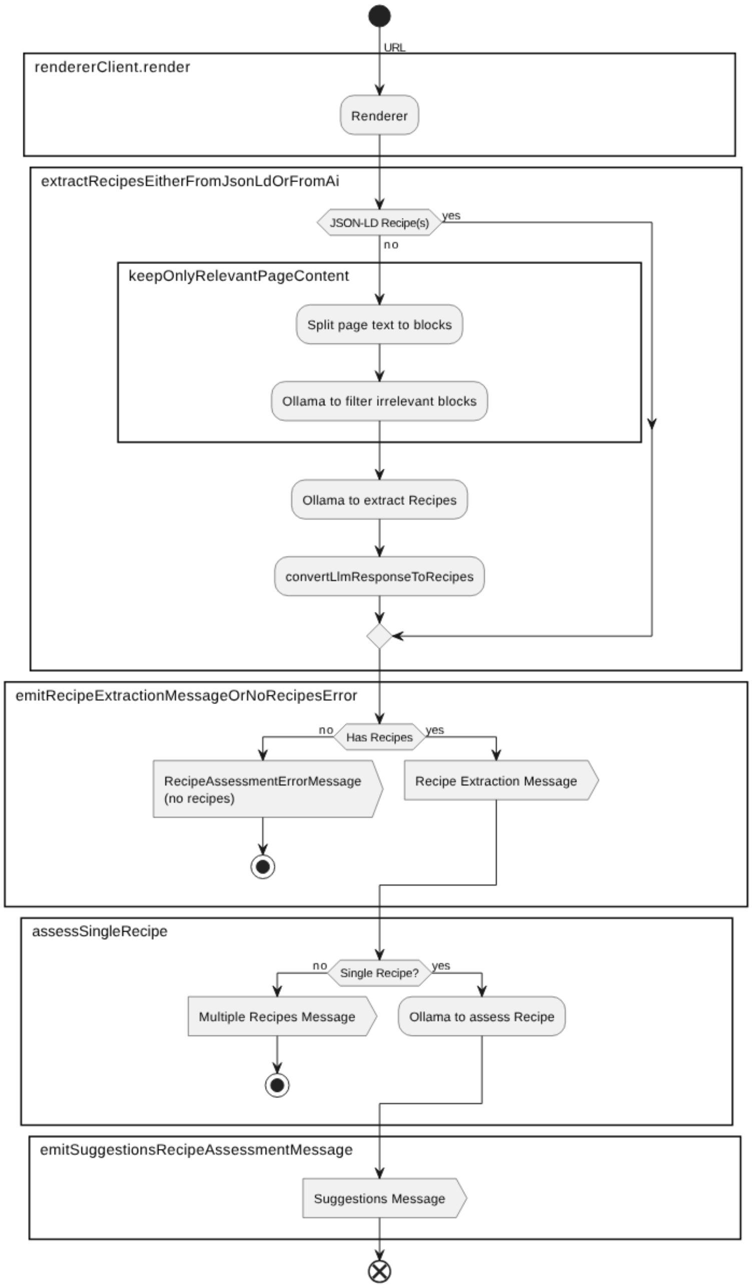

Figure 7: Flow diagram of recipe extraction process.

4.3.4. Technical Flow of AI Suggestions

As explained in the section 3.2, suggestions are limited to ingredient substitutions for both ICT solutions. The key distinction between RW and RCA lies in the personalization mechanism. While RW targets individual users, thus requires personalized suggestions, RCA is directed to influencers, whose content is visible to a broad and probably heterogenous audience. In this case, suggestions cannot be personalized, but rather aligned with general nutritional guidelines referring to an average person.

Dietary rules were defined by domain experts based on the most recent GBD Study. In the case of RW, substitutions are ranked according to the weight of the relevant GBD rule taking into account the personalization parameters (age and gender), as presented in Figure 8. Beyond rule ranking, the substitution logic also accounts for the functional role of each ingredient within the recipe, as well as any associated preparation or usage constraints (see Figure 9). The substitution logic can be phrased as follows:

Substitute ingredient X with Y, if the role of X is ROLE and RESTRICTION.

To illustrate the process more clearly, Figure 10 presents a concrete example. Suppose that the extracted recipe includes beef among its ingredients. Beef is treated as a trigger ingredient in the workflow and is associated with the rule "Decrease red meat". If the beef in the recipe is minced and used in a sauce, one suitable alternative to suggest would be cooked brown lentils, provided they are not intended for use in burgers without binders.

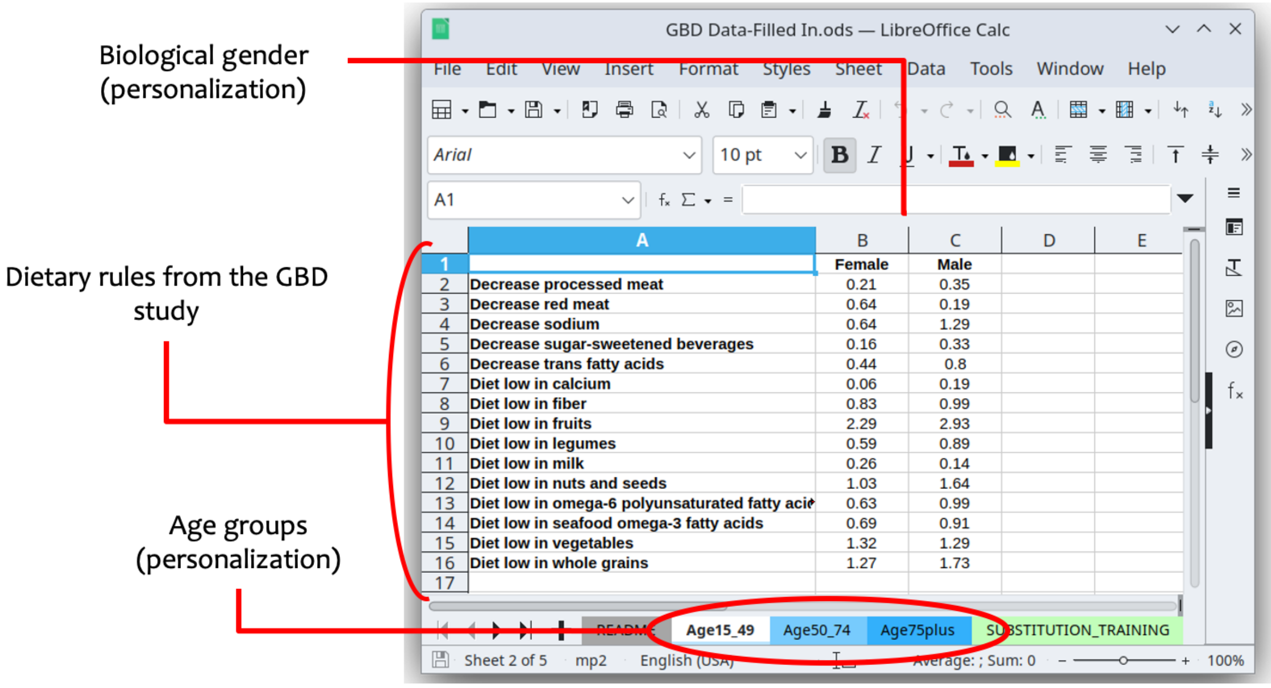

	A	B	C
1		Female	Male
2	Decrease processed meat	0.21	0.35
3	Decrease red meat	0.64	0.19
4	Decrease sodium	0.64	1.29
5	Decrease sugar-sweetened beverages	0.16	0.33
6	Decrease trans fatty acids	0.44	0.8
7	Diet low in calcium	0.06	0.19
8	Diet low in fiber	0.83	0.99
9	Diet low in fruits	2.29	2.93
10	Diet low in legumes	0.59	0.89
11	Diet low in milk	0.26	0.14
12	Diet low in nuts and seeds	1.03	1.64
13	Diet low in omega-6 polyunsaturated fatty aci	0.63	0.99
14	Diet low in seafood omega-3 fatty acids	0.69	0.91
15	Diet low in vegetables	1.32	1.29
16	Diet low in whole grains	1.27	1.73

Figure 8: Rules are ranked according to their contribution weights for three age groups and for male and female populations. Personalization within the system is guided by the gender and age parameters.

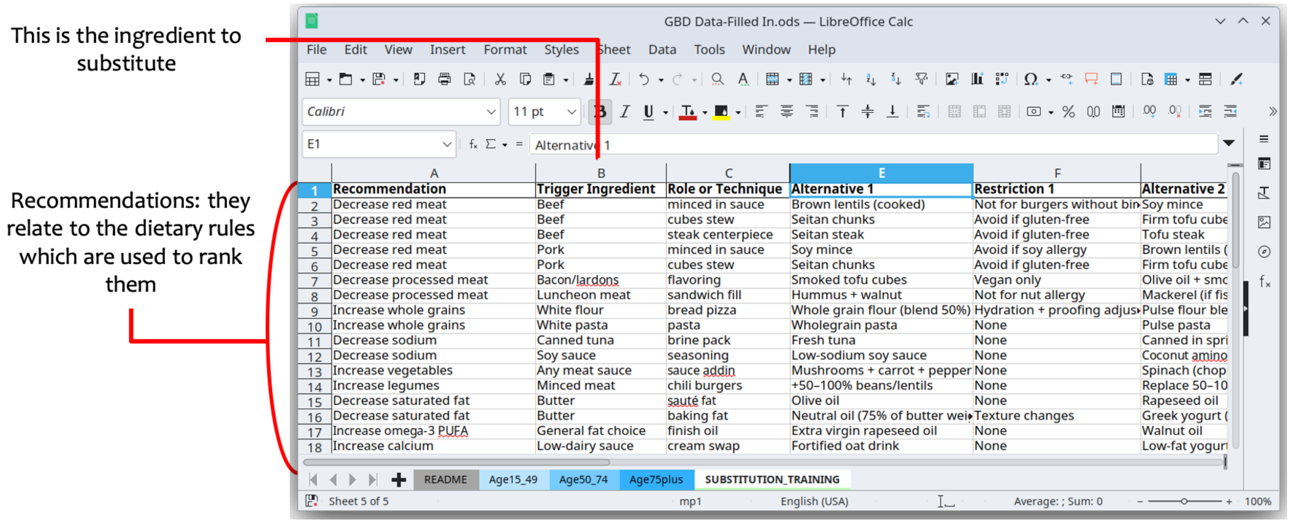

	A	B	C	E	F	
1	Recommendation	Trigger Ingredient	Role or Technique	Alternative 1	Restriction 1	Alternative 2
2	Decrease red meat	Beef	minced in sauce	Brown lentils (cooked)	Not for burgers without bin	Soy mince
3	Decrease red meat	Beef	cubes stew	Seitan chunks	Avoid if gluten-free	Firm tofu cube
4	Decrease red meat	Beef	steak centerpiece	Seitan steak	Avoid if gluten-free	Tofu steak
5	Decrease red meat	Pork	minced in sauce	Soy mince	Avoid if soy allergy	Brown lentils (
6	Decrease red meat	Pork	cubes stew	Seitan chunks	Avoid if gluten-free	Firm tofu cube
7	Decrease processed meat	Bacon/lardons	flavoring	Smoked tofu cubes	Vegan only	Olive oil + smo
8	Decrease processed meat	Luncheon meat	sandwich fill	Hummus + walnut	Not for nut allergy	Mackerel (if fis
9	Increase whole grains	White flour	bread pizza	Whole grain flour (blend 50%)	Hydration + proofing adjus	Pulse flour ble
10	Increase whole grains	White pasta	pasta	Wholegrain pasta	None	Pulse pasta
11	Decrease sodium	Canned tuna	brine pack	Fresh tuna	None	Canned in spri
12	Decrease sodium	Soy sauce	seasoning	Low-sodium soy sauce	None	Coconut amino
13	Increase vegetables	Any meat sauce	sauce addin	Mushrooms + carrot + pepper	None	Spinach (chop
14	Increase legumes	Minced meat	chili burgers	+50–100% beans/lentils	None	Replace 50–10
15	Decrease saturated fat	Butter	sauté fat	Olive oil	None	Rapeseed oil
16	Decrease saturated fat	Butter	baking fat	Neutral oil (75% of butter wei	Texture changes	Greek yogurt (
17	Increase omega-3 PUFA	General fat choice	finish oil	Extra virgin rapeseed oil	None	Walnut oil
18	Increase calcium	Low-dairy sauce	cream swap	Fortified oat drink	None	Low-fat yogurt

Figure 9: Create recommendations related to dietary factors and based on trigger ingredients to substitute.

Regarding RCA, the substitutions logic is very similar. The only difference is that the ranking cannot be materialized through the personalization factors. Therefore, we proceed with calculating the mean weights across both genders and all age groups and rank the dietary rules accordingly.

It is important to emphasize that the expert-defined mappings are not intended to be an exhaustive catalogue of possible alternatives. Rather, they serve as a background knowledge and conceptual guidance for the LLM, enabling it to generate suggestions that are aligned with an expert-approved rationale.

Given:

1. The extracted recipe,
2. User-specific information (only for RW),
3. The ranked dietary rules, and
4. Expert-defined mappings between trigger ingredients, substitutes, and constraints (Figure A 2),

a Large Language Model (LLM) is deployed to generate the final suggestions (see Figure 11). Dedicated prompt design is carried out for each use case (RW and RCA) to ensure appropriate contextualization. For the current experiments, nemotron-mini (4b) is deployed and integrated into the pipeline. It is important to note that although the same underlying model is selected across the system, it operates as distinct agents depending on the task definition and consequent prompt configuration. In earlier stages of the pipeline (section 4.3.3) the model functions as an internal agent, not exposed to the user, whose sole purpose is to reliably extract structured recipe information when JSON-LD is not reliable. During the suggestions workflow, the model is reconfigured through dedicated prompt engineering to operate as a nutritional advisory agent, whose outputs are presented to the end user, who may explicitly accept or decline the proposed substitutions.

The suggestions workflow itself is multistage and exploits the LLM for two conceptually different tasks: 1) identification of ingredient roles within the recipe, 2) generation of ingredient substitution suggestions. The overall workflow of the staged AI-suggestions process is presented in Figure 12. After recipe extraction (ingredients and steps), the top X ranked dietary rules/risks are retrieved. Note that in the case of RW rules are ranked based on personalization parameters, while in the case of RCA on population-level averages. The recipe ingredients are then categorized and matched against trigger ingredients associated with the ranked rules. The functional role of each matched ingredient is identified with respect to the specific recipe. We only retain the rows of the expert-defined mappings that correspond to matched trigger ingredients and top-ranked rules. Finally, the LLM is prompted with the filtered expert mappings as background knowledge and guidance, rather than a fixed lookup table. The model is allowed to generate substitutions that are either directly listed in the expert mappings or inspired by them, provided that they adhere to dietary rules, ingredient roles and restriction constraints. At present, the LLM-related stages of suggestions rely on ongoing prompt development and refinement, with different prompts designed for role identification and for substitution generation. This separation reflects the fundamentally different reasoning requirements of the two tasks and allows experimentation and iterative improvement without hard-coding assumptions into the system.

This staged and guided approach minimizes the amount of information passed to the LLM at inference time, thereby preserving the available context window and reducing the risk of hallucinations. At the same time, it enables the generation of nutritionally sound, expert-informed suggestions while retaining sufficient flexibility to accommodate diverse recipes and cooking practices.

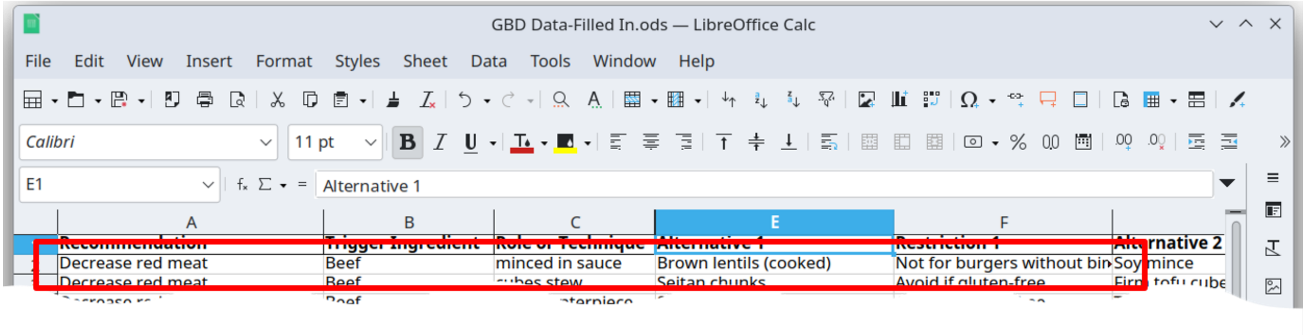

Figure 10: Example of the substitution logic accounting for trigger ingredient and its associated rule/risk, its functional role in the recipe and the suggested alternatives together with the corresponding restrictions.

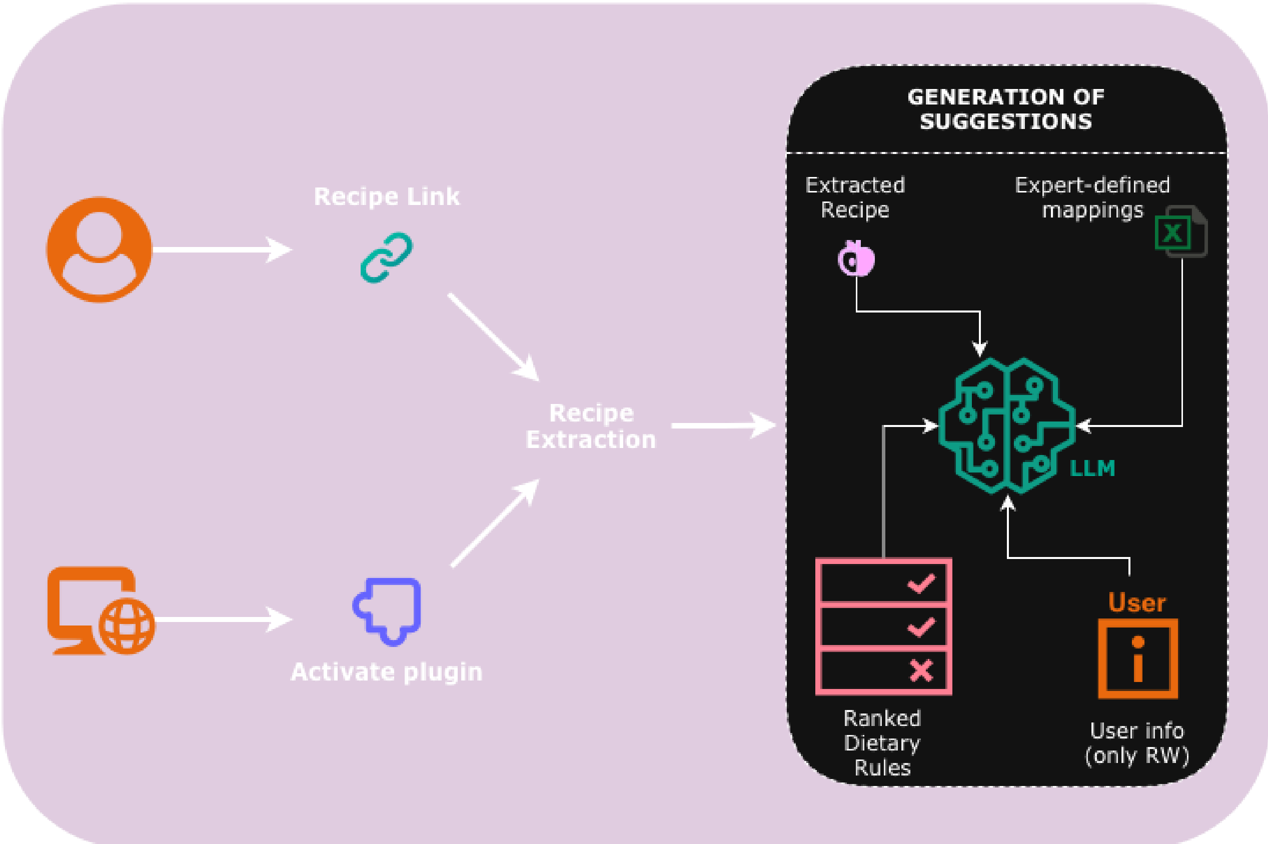

Figure 11: End-to-end workflow overview of AI suggestions. A recipe is provided either via a direct link (RW) or through plugin activation (RCA) and is subsequently processed by the recipe extraction component. The extracted recipe, together with the ranked dietary rules (personalized for RW or population-based for RCA), expert-defined substitution mappings, and user information (only in RW), is passed to a Large Language Model (LLM), which generates nutritionally informed ingredient substitution suggestions.

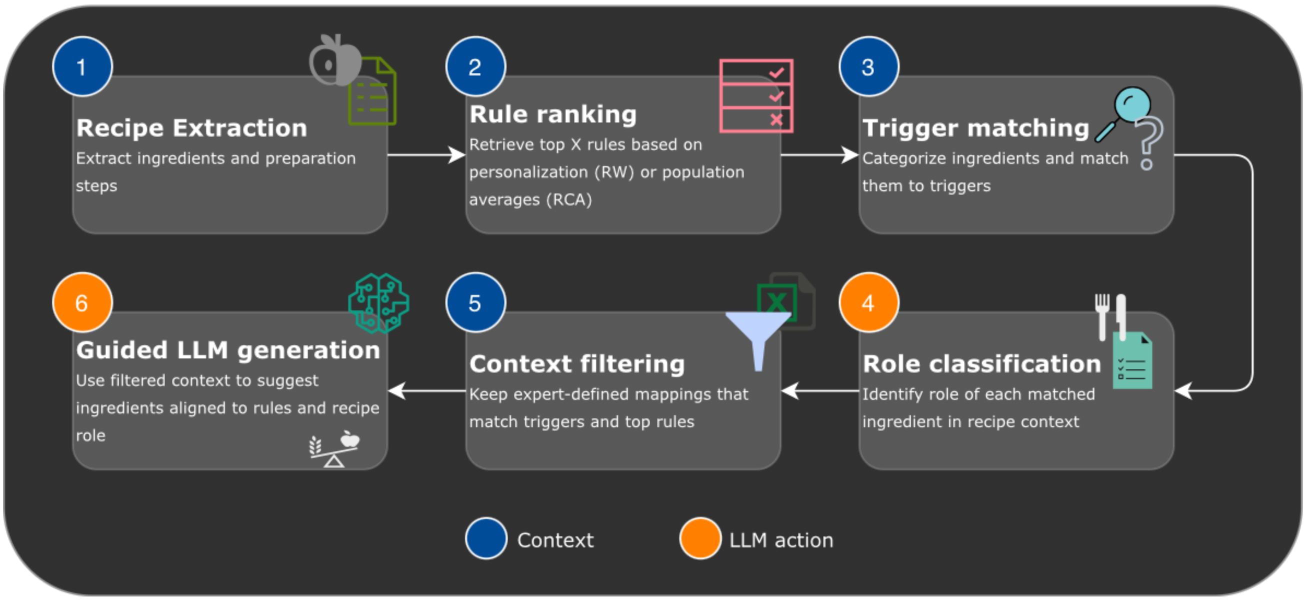

Figure 12: Zoom into generation of suggestions: Following recipe extraction, dietary rules are ranked based on personalization parameters (RW) or population-level averages (RCA). Recipe ingredients are then matched to trigger ingredients associated with the top-ranked rules, and their functional role within the recipe is identified. Expert-defined substitution mappings are filtered to retain only those relevant to the matched triggers and rules. Finally, the LLM is prompted with this filtered context to generate guided, expert-aligned ingredient substitution suggestions while minimizing hallucinations and preserving contextual relevance.

4.4. TECHNICAL RESTRICTIONS

This section identifies the main technical restrictions foreseen for the development and operation of the two ICT solutions, namely RecipeWatch (RW) and Responsible Cooking Alliance (RCA), and outlines the mitigation strategies planned to address them.

A first restriction is associated with RecipeWatch (RW). Unlike browser extensions, mobile applications cannot directly access the content of external web pages. As a result, RW relies on server-side rendering and extraction mechanisms to retrieve recipe content, as explained in the section 4.3.2. This approach may fail in certain scenarios, including:

- Websites that require user authentication before displaying the recipe,
- Websites that enforce human verification mechanisms (e.g. CAPTCHA), and
- Websites that actively block automated or headless browsers.

Although significant effort has been invested in configuring the controlled browser environment to closely resemble a standard user browser, it cannot be guaranteed that all websites will behave consistently. Given the large number of existing recipe websites, exhaustive testing across all possible cases is not feasible. As a mitigation measure, the system is designed to gracefully handle such failures by detecting inaccessible pages and informing the user accordingly. Continuous monitoring and incremental improvements to the rendering strategy will be applied as new incompatibilities are identified during operation.

A further technical restriction applies to the extraction of recipes from third-party websites, which is a fundamental step, since the quality of the extracted recipes directly impacts all subsequent processing stages. As outlined in the section 4.3.3 our hybrid design aims to overcome the inherent challenges derived from the heterogeneity of web environments.

An additional restriction concerns the accuracy and robustness of the AI-based logic used to identify ingredients, group them appropriately, infer their functional role in the recipe and successfully apply the substitutions logic in a meaningful and context-aware manner. This task is inherently complex, as ingredients may serve different purposes depending on context and recipes often rely on implicit culinary knowledge that is not stated in the text. Errors at this stage could result in inappropriate or impractical substitutions, thereby affecting the quality and usefulness of the recommendations. To mitigate this complexity, the suggestions workflow has been intentionally designed as a multistage process that decomposes the overall task into smaller, more concrete sub-tasks. This approach allows each stage to focus on a well-defined objective, thereby minimizing the length and complexity of the input provided to the LLM, improving prompt efficiency, and reducing the risk of reasoning errors. In parallel, expert-defined mappings and dietary rules are selectively filtered and introduced only when relevant, ensuring that domain knowledge is exploited in a targeted and effective manner rather than as a lengthy input.

Overall, the identified technical restrictions are addressed through targeted architectural choices and iterative refinement strategies. Concerning the AI components, the restrictions identified reflect the inherent challenges of deploying AI-driven functionalities, while ensuring reliable results and controlled reasoning. By explicitly accounting for these constraints in the system architecture and development process, the project aims to deliver reliable, meaningful, and user-trustworthy recommendations. For ease of reference, Table X summarizes the main technical restrictions identified for RW and RCA, together with the mitigation strategies adopted.

Solution	Technical restriction	Mitigation strategy
RW	Limited access to third-party recipe websites due to authentication, CAPTCHA, or blocking of automated browsers	Server-side rendering with controlled browser configuration; graceful failure detection and user notification; continuous monitoring and refinement
RW	Heterogeneity and variability of recipe website structures affecting extraction accuracy	Hybrid recipe extraction approach; iterative tuning based on observed edge cases
RW & RCA	Complexity of AI-based ingredient identification and substitution logic	Multistage suggestion workflow; reduced prompt complexity; selective use of expert-defined rules
RW & RCA	Inherent variability and limitations of AI-driven reasoning	Controlled reasoning workflows; architectural safeguards; iterative refinement during operation

Table 2: Summary of technical restrictions and mitigation strategies.

5. USABILITY TESTING

Usability testing constitutes a key quality assurance and validation activity within the DietWise project, ensuring that the developed ICT solutions are usable, understandable, and aligned with the needs and expectations of their target user groups. In accordance with the Grant Agreement, usability testing will be performed both on application prototypes and on coded beta versions. At the time of writing, the prototypes of the two user-facing applications -Recipe Watch (RW) and Responsible Cooking Alliance (RCA)- are under development. Therefore, this section describes the planned usability testing strategy, detailing the methodology, tools, evaluation criteria, and feedback integration process that will be applied once the prototypes and beta versions become available in M18 as planned in the GA.

5.1. GOALS AND SCOPE

The overarching goal of the usability testing activities is to verify that the DietWise applications:

- Support users in completing their core tasks effectively and efficiently.
- Provide clear, transparent, and trustworthy feedback and recommendations.
- Minimize cognitive load and user errors.
- Are accessible to users with varying levels of digital literacy.
- Foster engagement and acceptance of AI-assisted decision support.

Given the different nature and target audiences of the two applications, usability testing will explicitly account for the distinct usage contexts of RecipeWatch (citizens, mobile-first usage) and Responsible Cooking Alliance (influencers, desktop browser-based usage).

Target User Groups

Usability testing will involve representative users from the intended audiences of each application:

- **Recipe Watch (RW):**
  - Citizens interested in nutrition and sustainability
  - Users with different age groups and levels of familiarity with mobile applications
  - Both registered and anonymous usage scenarios where applicable
- **Responsible Cooking Alliance (RCA):**
  - Food influencers and content creators
  - Users accustomed to managing web-based content
  - Users with prior experience publishing recipes online

Participant recruitment will aim to capture diversity in experience, expectations, and usage patterns, also keeping in mind the three defined stakeholder groups of the co-creation activities, namely citizens, vulnerable citizens and influencers, in order to identify usability issues that may not emerge from homogeneous user samples.

5.2. USABILITY TESTING PHASES

Usability testing will be conducted in two complementary phases, corresponding to increasing levels of implementation maturity.

Phase 1: Prototype-Based Usability Testing

Once interactive prototypes (low- and high-fidelity) become available, usability testing will be conducted. This phase focuses on early validation of interaction design and information architecture, before full implementation. It aims to identify conceptual misunderstandings, unclear terminology, navigation issues, and mismatches between user expectations and system behavior.

Methods foreseen in this phase include:

- Task-based usability testing (moderated and/or unmoderated)
- Scenario-driven interaction walkthroughs
- First-click and navigation clarity tests
- Post-task and post-session questionnaires

Indicative usability scenarios for the RecipeWatch:

- Discovering the main functionality of the application
- Submitting a recipe URL and understanding the extracted recipe
- Reviewing and interpreting ingredient substitution suggestions
- Accepting or rejecting substitutions and understanding their impact
- Understanding aggregated feedback from other users
- Identifying the final modified recipe and the applied changes

Indicative usability scenarios for the Responsible Cooking Alliance:

- Activating the browser extension on a recipe page
- Locating and opening the RCA interface within the browser
- Understanding the structure and meaning of the feedback provided
- Distinguishing between general advice and ingredient-specific suggestions
- Interpreting recipe rankings and sustainability indicators

Phase 2: Beta Testing of Coded Applications

Following the implementation of functional application versions, usability testing will be extended to beta testing in real usage environments.

- RecipeWatch will be distributed through the usual distribution channels of the app stores for each technology – iOS/Android.
- Responsible Cooking Alliance will be distributed via controlled browser extension testing channels.

This phase enables evaluation of usability under realistic conditions, including:

- Performance and responsiveness
- Latency related to AI-based processing
- Stability across devices and browsers
- Integration of usability with technical constraints

Data collection methods in this phase include:

- In-app usability questionnaires
- Voluntary user feedback forms
- Observation of task completion and abandonment patterns
- Qualitative feedback on clarity, trust, and perceived usefulness

Concerning usability metrics and evaluation criteria, across both phases, usability evaluation will rely on a combination of qualitative and quantitative indicators, including:

- Task completion rates
- Time required to complete key tasks
- Frequency and type of user errors
- Perceived ease of use and clarity
- User confidence in and trust toward recommendations
- Overall satisfaction and intention to reuse

These metrics will allow comparison across iterations and identification of high-impact usability issues. Additionally, all usability testing outcomes will be systematically documented and analyzed. Identified issues will be categorized according to severity (e.g., critical, major, minor) and mapped to specific application components or interaction flows. Findings will directly inform:

- User interface refinements
- Interaction flow adjustments
- Clarification of labels, explanations, and visual cues
- Improvements in transparency around AI-generated suggestions

This iterative, user-centered refinement process ensures that usability considerations are continuously integrated into the technical development lifecycle of DietWise. By planning usability evaluation activities across multiple development stages, the project ensures early risk mitigation, improved user acceptance, and higher overall quality of the delivered ICT solutions.

6. CONCLUSIONS

This deliverable (D4.2) consolidates the outcomes of WP4 co-development activities and documents how stakeholder insights and expert input have been translated into the initial design and technical implementation of the DietWise ICT solutions: Recipe Watch (RW) for citizens and Responsible Cooking Alliance (RCA) for influencers. It provides a traceable link between the requirements gathered through workshops and upstream evidence (WP2 and WP3) and the resulting functional priorities, workflows, and interface expectations that are now being implemented.

D4.2 further establishes the technical grounding required for the next project phases by presenting the shared system architecture underpinning both applications, including identity management, backend orchestration, recipe rendering and extraction mechanisms, and the integration of AI components for recipe assessment and substitution generation. In parallel, the deliverable identifies key technical and adoption-related restrictions, most notably recipe extraction variability, constraints of server-side rendering for RW, and the robustness and acceptability of AI-driven substitutions, and outlines mitigation strategies aligned with an iterative development approach. These measures aim to safeguard system reliability, transparency, and user trust as development progresses.

Finally, the deliverable defines a structured usability testing strategy to be applied once interactive prototypes and coded beta versions become available. By combining prototype-based testing with real-world beta testing, DietWise will ensure that both RW and RCA are validated with representative users, enabling evidence-based refinement of user experience, trust mechanisms, and the overall effectiveness of recommendations. In this way, D4.2 provides a robust foundation for subsequent pilot preparation and deployment activities (WP5 and WP7), supporting the project's progression from co-created concepts to implementable, testable, and scalable ICT solutions.

Annex A

	Components	Limited/Encouraged	Component in Recipe	Point/Component		Original Score	Updated Score
1–2	Components	Limited/Encouraged	Component in Recipe	Point/Component		Original Score	Updated Score
3	processed meat	Limited	Yes/No	1/0		0	?
4	red meat	Limited					
5	sodium	Limited					
6	sugar-sweetened beverages	Limited					
7	trans fatty acids	Limited					
8	calcium	Encouraged					
9	fiber	Encouraged					
10	fruits	Encouraged					
11	legumes	Encouraged					
12	milk	Encouraged					
13	nuts and seeds	Encouraged					
14	omega-6 polyunsaturated fatty acids	Encouraged					
15	seafood omega-3 fatty acids	Encouraged					
16	vegetables	Encouraged					
17	whole grains	Encouraged					
18				0			
19	*The score can be presented **for the original and improved** (based on the AI recommendations recipes) to encourage or emphasize how the proposed changes enable healthier choices*						
20	*Next step is testing a number of recipes to understand limitations, and "degree" of changes to the score after accepting recommendations*						

Figure A 1: The proposed scoring logic, uses the GBD dietary risk factors as the set of components to rate how well a recipe aligns with healthy eating. Each component corresponds to a dietary factor that should either be limited or encouraged. By assessing the presence or absence of these factors in a given recipe, the system computes a composite score. The higher the score, the more the recipe contributes to a healthy diet (and the fewer risk factors it poses).

	A	B	C	D	E	F	G	H	I	J	K	L
1	**Recommendation**	**Trigger Ingredient**	**Role or Technique**	**Cuisine Context**	**Alternative 1**	**Restriction 1**	**Alternative 2**	**Restriction 2**	**Alternative 3**	**Restriction 3**	**Equivalence Notes**	**Technique Notes**
2	Decrease red meat	Beef	minced in sauce	Italian	Brown lentils (cooked)	Not for burgers without binder	Soy mince	Avoid if soy allergy	Mushroom–lentil mix	Use only in sauces	1:1 (200g beef → 200g cooked lentils) or 200g beef → 120g lentils + 150g mushrooms	Finely chop mushrooms; dry sauté before adding
3	Decrease red meat	Beef	cubes stew	Any	Seitan chunks	Avoid if gluten-free	Firm tofu cubes	Not suitable for quick stir-fry	Young jackfruit chunks	Lower protein content	1:1 by weight	Simmer longer for seitan; coat tofu in starch before braise
4	Decrease red meat	Beef	steak centerpiece	Any	Seitan steak	Avoid if gluten-free	Tofu steak	Press 20 min before searing	Grilled portobello	Lower protein/iron	1:1 piece count (approx. 150–200g)	High-heat sear for crust
5	Decrease red meat	Pork	minced in sauce	Any	Soy mince	Avoid if soy allergy	Brown lentils (cooked)	Texture softer	Mushroom–lentil mix	Use only in sauces	1:1 by weight	Pre-brown
6	Decrease red meat	Pork	cubes stew	Any	Seitan chunks	Avoid if gluten-free	Firm tofu cubes	Press first	Jackfruit chunks	Lower protein	1:1 by weight	Simmer gently
7	Decrease processed meat	Bacon/lardons	flavoring	Any	Smoked tofu cubes	Vegan only	Olive oil + smoked paprika	Not suitable for vegan if anchovy used	Mushroom + soy sauce	Use small amounts	Replace 50–100g bacon with 120g smoked tofu	Render tofu to crisp
8	Decrease processed meat	Luncheon meat	sandwich fill	Any	Hummus + walnut	Not for nut allergy	Mackerel (if fish allowed)	Not vegan	Roasted chickpea spread	Lower protein than meat	1 slice → 40–60g spread	Spread on wholegrain bread
9	Increase whole grains	White flour	bread pizza	Italian	Whole grain flour (blend 50%)	Hydration + proofing adjust	Pulse flour blend (20–30%)	Texture changes	Spelt (partial)	Gluten content differs	Start 30–50% blend	Increase hydration +5–10%
10	Increase whole grains	White pasta	pasta	Italian	Wholegrain pasta	None	Pulse pasta	None	50/50 mix (white+wholegrain)	None	1:1 weight	Track cook time (firmer bite)
11	Decrease sodium	Canned tuna	brine pack	Any	Fresh tuna	None	Canned in spring water	Rinse before use	Half tuna + half beans	Texture mix	1:1 weight	Rinse canned tuna to reduce salt
12	Decrease sodium	Soy sauce	seasoning	Asian stirfry	Low-sodium soy sauce	None	Coconut aminos	Sweeter profile	Miso paste + water	Not 1:1 in stir-fries	Use 50–75% of original sodium	Add gradually
13	Increase vegetables	Any meat sauce	sauce addin	Italian	Mushrooms + carrot + pepper	None	Spinach (chopped)	Water down risk	Courgette (grated)	Watch dilution	+150–250g extra veg per 4p	sweat veg before adding
14	Increase legumes	Minced meat	chili burgers	Any	+50–100% beans/lentils	None	Replace 50–100% with soy mince	Check soy allergy	Lentil–oat burger mix	Needs binder	1:1 replacement or add 80–150g beans per 200g meat	Add binder (oats/egg/flax)
15	Decrease saturated fat	Butter	sauté fat	Any	Olive oil	None	Rapeseed oil	None	Half butter/half oil	Still some sat fat	1:1 by volume for sauté	Heat oil first
16	Decrease saturated fat	Butter	baking fat	Any	Neutral oil (75% of butter weight)	Texture changes	Greek yogurt (30–50% swap)	Moisture shift	Apple sauce (up to 30%)	Flavor impact	100g butter → 75g oil; yogurt swaps vary	Adjust liquids and baking time
17	Increase omega-3 PUFA	General fat choice	finish oil	Any	Extra virgin rapeseed oil	None	Walnut oil	Not for nut allergy	Flaxseed oil	Use cold only	Use 1–2 tbsp per 4p	Use as finishing oil only
18	Increase calcium	Low-dairy sauce	cream swap	Any	Fortified oat drink	None	Low-fat yogurt	Not vegan	Silken tofu cream	Vegan only	1:1 for cream (watch heat)	Blend smooth; avoid boiling yogurt
19	Increase fiber	White rice	staple	Any	Brown rice	None	Half brown/half white	None	Add barley (soup/stews)	Gluten content	1:1 weight	Soak brown rice or cook longer
20	Increase fiber	Refined bread	bread	Any	Wholegrain bread	None	Sourdough wholegrain	None	Add seeds mix	Allergy seeds	1:1 weight	Hydration adjust minimal
21	Decrease trans fatty acids	Margarine (non-HO)	baking fat	Any	Butter or HO margarine	Check sat fat	Olive oil swap (partial)	Texture changes	Avocado purée (partial)	Flavor shift	Case-by-case	Adjust bake time
22	Increase nuts and seeds	Salad topping	topping	Any	Sunflower seeds	None	Pumpkin seeds	None	Walnuts	Not for nut allergy	Add 20–40g per 4p	Toast lightly
23	Increase seafood omega-3	Protein choice	swap in	Any	Mackerel	None	Sardines	None	Salmon	None	1:1 weight	Grill/bake
24	Decrease sugar-sweetened beverages	SSB	beverage	Any	Water + fruit infusion	None	Unsweetened tea	None	Sparkling water	None	N/A	N/A
25	Increase milk	Low-dairy breakfast	swap in	Any	Low-fat milk	None	Fortified plant drink	None	Skyr/yogurt	Not vegan	1:1 volume	Cold prep
26	Increase legumes	Roerbak proteïne	stirfry protein	Asian stirfry	Tempeh	None	Firm tofu	None	Seitan	Not gluten-free	1:1 weight	Press tofu; wok heet
27	Increase vegetables	Pasta dishes	veg boost	Italian	Add 150–250g veg per 4p	None	Swap 50% pasta → veg noodles	None	Add legumes to sauce	None	See note	Blanch or sauté veg noodles
28	Decrease sodium	Stock cube	broth base	Any	Low-salt stock	None	Homemade veg stock	None	Half cube + extra herbs	None	Use 50% sodium	Simmer with aromatics
29	Decrease red meat	Lamb	curry cubes	Indian	Chickpea + aubergine mix	Lower protein	Paneer (if dairy ok)	Not vegan	Tofu cubes	None	1:1 by volume	Roast aubergine first
30	Decrease saturated fat	Cream	sauce enricher	European	Evaporated milk	Not vegan	Skimmed milk + roux	Thinner body	Cashew cream	Not for nut allergy	200ml cream → 200ml evap or 150ml skim + roux	Simmer gently
31	Increase whole grains	White flour	roux binder	Any	Wholemeal flour (25–50%)	May darken sauce	Oat flour (partial)	Gluten different	Chickpea flour (partial)	Flavor shift	Replace 25–50%	Cook flour well
32	Increase fiber	White couscous	staple	Middle Eastern	Wholewheat couscous	None	Bulgur	None	Quinoa	None	1:1 volume	Hydration per pack

Figure A 2: The parameters used in the rules-based system are derived from the Global Burden of Disease Study (2023): https://ghdx.healthdata.org/gbd-2023. Dietary risk factors are listed according to their contribution to total cause-specific mortality. Based on each parameter/dietary risk, experts developed a dataset of relevant recipe optimizations, as knowledge pool for the LLM recommendations. The dataset works off the basis of a trigger ingredient in the recipe and provides alternatives for cooking technique(s) and alternative ingredient(s)

	A	B	C
1		**Female**	**Male**
2	**Diet low in whole grains**	1.27	1.73
3	**Diet low in vegetables**	1.32	1.29
4	**Diet low in seafood omega-3 fatty acids**	0.69	0.91
5	**Diet low in omega-6 polyunsaturated fatty acid**	0.63	0.99
6	**Diet low in nuts and seeds**	1.03	1.64
7	**Diet low in milk**	0.26	0.14
8	**Diet low in legumes**	0.59	0.89
9	**Diet low in fruits**	2.29	2.93
10	**Diet low in fiber**	0.83	0.99
11	**Diet low in calcium**	0.06	0.19
12	**Decrease trans fatty acids**	0.44	0.8
13	**Decrease sugar-sweetened beverages**	0.16	0.33
14	**Decrease sodium**	0.64	1.29
15	**Decrease red meat**	0.64	0.19
16	**Decrease processed meat**	0.21	0.35

Figure A 3: The relative contribution of each GBD factor for age group 15-49, for male and female populations.

	A	B	C
1		**Female**	**Male**
2	**Decrease processed meat**	0.56	0.75
3	**Decrease red meat**	0.89	0.73
4	**Decrease sodium**	2.72	4.48
5	**Decrease sugar-sweetened beverages**	0.29	0.42
6	**Decrease trans fatty acids**	0.76	0.99
7	**Diet low in calcium**	0.3	0.09
8	**Diet low in fiber**	0.9	0.9
9	**Diet low in fruits**	3.68	4.16
10	**Diet low in legumes**	1.08	1.32
11	**Diet low in milk**	0.49	0.27
12	**Diet low in nuts and seeds**	1.67	2.07
13	**Diet low in omega-6 polyunsaturated fatty acid**	1.04	1.34
14	**Diet low in seafood omega-3 fatty acids**	1.09	1.07
15	**Diet low in vegetables**	1.96	1.89
16	**Diet low in whole grains**	2.42	2.78

Figure A 4: The relative contribution of each GBD factor for age group 50-74, for male and female populations.

	A	B	C
1		Female	Male
2	Decrease processed meat	0.52	0.59
3	Decrease red meat	0.61	0.71
4	Decrease sodium	2.59	3.73
5	Decrease sugar-sweetened beverages	0.21	0.27
6	Decrease trans fatty acids	0.55	0.6
7	Diet low in calcium	0.21	0.05
8	Diet low in fiber	0.79	0.68
9	Diet low in fruits	2.68	2.88
10	Diet low in legumes	1.12	1.02
11	Diet low in milk	0.34	0.17
12	Diet low in nuts and seeds	1.47	1.41
13	Diet low in omega-6 polyunsaturated fatty acid:	0.75	0.8
14	Diet low in seafood omega-3 fatty acids	0.98	0.8
15	Diet low in vegetables	1.76	1.55
16	Diet low in whole grains	2.02	1.97

Figure A 5: The relative contribution of each GBD factor for age group 75plus, for male and female populations.