D4.5 First prototype of the predictive dashboard

WP4: Decision-Making Tools for Nutritional and Sustainable Food Systems

Responsible Author: Hans Houf (Jibe)

Funded by the European Union

Document information

Programme and call	Horizon Europe: HORIZON-CL6-2023-GOVERNANCE-01
Project acronym	SOSFOOD
Project name	Sustainability Optimization for Secure Food Systems
Start of the project	February 2024
Duration	48 months
Project coordination	Jesús Simal Gándara (UVigo) Eloy Miranda Garcia (UVigo)

Deliverable title	First prototype of the predictive dashboard
Deliverable number	4.5
Nature of the deliverable	Report
Dissemination level	Public
Work Package	WP4 – Decision-Making Tools for Nutritional and Sustainable Food Systems
Leading Beneficiary	Jibe
Delivery date	M12 (January 2025)

Funded by the European Union

Document History

Time	Session	Date	Authors
01	First draft	7/01/2025	Carolina Gogova, Cosmin Ularu
02	Second draft	23/01/2025	Hans Houf Cosmin Ularu Carolina Gogova

Funded by the European Union

Table of Contents

Funded by the
European Union

Executive Summary

This document presents the development, functionality, and testing process of the SOSFood Predictive Dashboard, a critical tool within the project's mission to transform food systems through data-driven innovation. Designed to provide stakeholders with actionable insights, the dashboard integrates advanced predictive capabilities powered by AI, enabling industries, policymakers, and researchers to make informed decisions grounded in robust data and methodologies.

The dashboard's key features include its multiplatform accessibility, customizable data visualizations, and role-based access system, making it adaptable to the needs of a diverse range of users. With dynamic views for graphs, maps, and tables, users can analyze data across indicators such as climate, soil, economic variables, and crop-related metrics. The dashboard's ability to process historical data, current conditions, and future predictions, complete with uncertainty measures, ensures a comprehensive perspective on sustainability trends and challenges.

The technical foundation of the dashboard is built on Kibana, chosen for its capability to handle large datasets, create interactive visualizations, and seamlessly integrate with the project's data infrastructure. Its deployment, supported by Kubernetes and Elasticsearch, ensures a secure, scalable, and accessible system that aligns with the technical requirements outlined in the grant agreement.

A rigorous testing and validation process ensures the dashboard meets the highest standards of usability, accuracy, and relevance. This process includes both controlled environment testing, conducted under WP4, and real-world validation through case studies in Galicia, Athens, and Lithuania under WP5. Feedback loops from these activities drive iterative improvements, ensuring the dashboard evolves to meet user needs and remains adaptable to various operational contexts.

By synthesizing contributions from WP1, WP2, WP3, and WP5, the dashboard serves as a cornerstone of the SOSFood project's efforts to promote sustainability, transparency, and collaboration across food systems. With its first prototype already deployed and future refinements planned, the dashboard demonstrates the project's commitment to providing practical, impactful solutions for addressing complex food system challenges.

This executive summary provides an overview of the document's content, including the technical specifications, main features, and validation strategies of the predictive dashboard. It sets the stage for a detailed exploration of how the dashboard contributes to the SOSFood project's vision of creating resilient and sustainable food systems.

Funded by the
European Union

1. Introduction

The SOSFood project is dedicated to advancing sustainability in food systems by leveraging data-driven insights, innovative technologies, and multi-stakeholder collaboration. As part of this mission, the **Predictive Dashboard** has been developed as a central tool to analyze, visualize, and interpret complex datasets, enabling stakeholders to make informed, evidence-based decisions. By integrating predictive capabilities powered by AI, the dashboard offers a user-friendly platform for visualizing trends, exploring key sustainability indicators, and understanding the impacts of different variables across time and regions.

The dashboard is designed to address the needs of diverse users, from policymakers to industry leaders and researchers. Its primary purpose is to empower decision-makers to act on solid, scientifically-backed predictions, while also providing transparency into the data and methodologies used. This not only ensures usability but also builds trust among stakeholders, fostering greater alignment toward shared sustainability goals.

By synthesizing data from multiple work packages, the dashboard represents the collective efforts of the SOSFood consortium. It incorporates research on datasets from WP1, sustainability metrics and methodologies from WP2, AI-driven predictive models from WP3, controlled environment testing from WP4 and validation feedback from WP5. This integration of interdisciplinary expertise ensures that the dashboard is robust, adaptable, and aligned with the project's objectives to drive meaningful change in food systems.

This document outlines the development, features, and testing of the predictive dashboard, along with its technical requirements and real-world applications. Furthermore, it describes its design, functionality, and validation process, as well as the next steps for its refinement and deployment.

2. Technical requirements

The technical requirements for the predictive dashboard were determined based on the needs of the intended users, the nature of the input data, and the methods for visualizing that data effectively. These considerations guided the selection of technologies and specifications to ensure that the dashboard could meet the project's objectives while remaining functional and user-friendly.

The use of Kibana, an open-source data visualization tool, was specified in the Grant Agreement and chosen for its compatibility with the dashboard's requirements, including its ability to manage large datasets and create dynamic visualizations across multiple platforms. Kibana offers advanced visualization tools for large and complex datasets as well as built-in role management and customizable dashboards to adapt to different user needs which make it a suitable choice to fulfil the overarching goals, which are:

Funded by the European Union

- To provide stakeholders with actionable insights to improve sustainability in food systems
- To ensure the tool is accessible, adaptable, and transparent for all users
- To create a system that integrates diverse datasets and predictive insights into a unified platform

2.1 Main features of the dashboard

1. A multiplatform user-friendly view (available on both mobile and desktop)

2. Data visualisation:
    a. Creation of charts, graphs, and maps to visualize predictions across parameters such as climate, soil, economic indicators, and crop-related data.
    b. A menu listing all available parameters, allowing users to navigate and explore datasets from multiple perspectives.

3. User roles:
    a. Guest Users: Can view dashboards and visualizations but are restricted from editing or deleting any content.
    b. Dashboard Designers: Have permissions to create, edit, and manage dashboards and visualizations within their scope.
    c. Admin Users: Manage system configurations and user permissions.
    d. Additional user roles can be configured if needed.

4. Time and Data Filters:
    Users can adjust time granularity (e.g., daily, weekly, monthly) and filter data by specific parameters or regions to suit their needs.

5. View types
    Different views options such as graph view, map view, and table view

Some of these basic requirements can be seen in the mock ups below.

Funded by the European Union

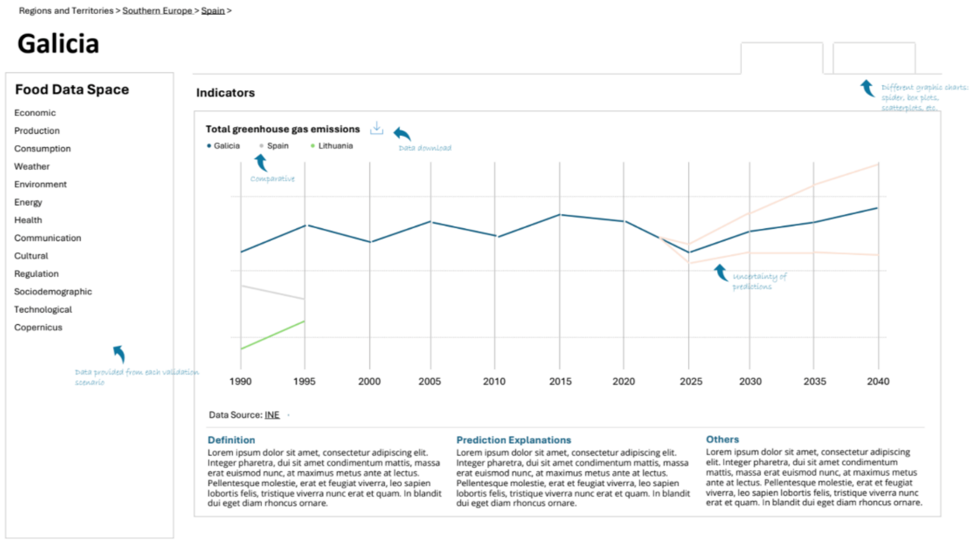

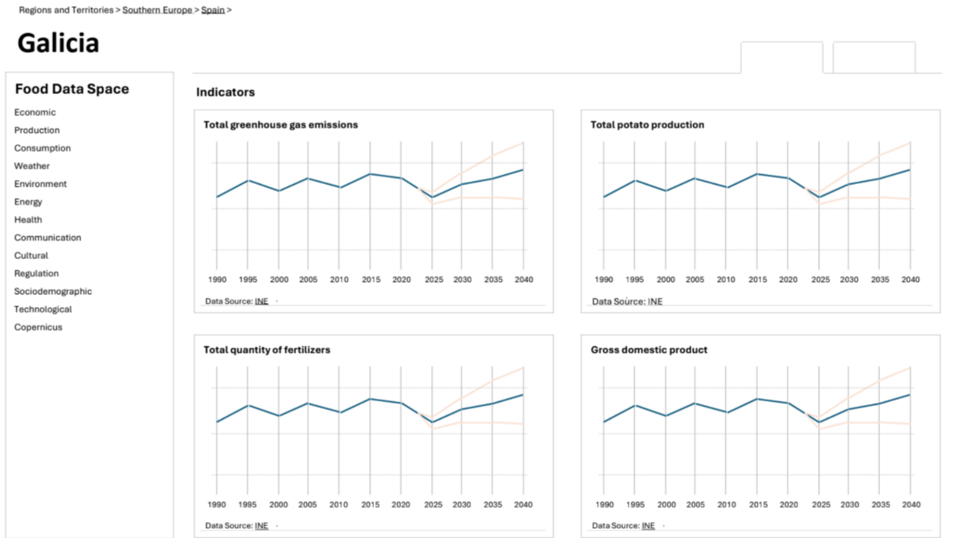

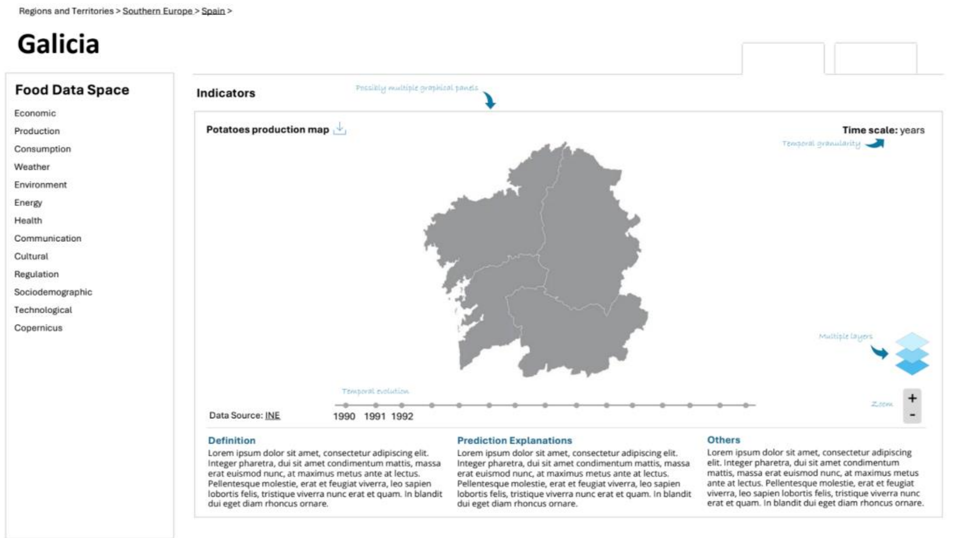

Funded by the European Union

2.2 Input for the Predictive Dashboard

The dashboard processes two primary types of input data:

1. **Single Predictions**: Represent specific data points for individual parameters.
2. **Time Series Predictions**:
   - Streams of data over a defined time range, including:
   - Historical data showing past trends.
   - Current data representing real-time conditions.
   - Forecasted data predicting future trends with associated uncertainty measures.

These inputs allow users to gain a holistic view of parameter evolution over time, supported by graphical representations of confidence levels and explanations of predictive methodologies.

By enabling both broad exploration and detailed analysis, the dashboard supports a wide range of users—from policymakers needing big-picture insights to researchers looking for granular details. The role-based access system keeps things simple and secure, while the flexibility in visualizing and filtering data ensures the dashboard remains adaptable to different contexts.

3. First prototype

The first prototype of the dashboard includes three main views for analyzing indicators: **Graph View**, **Map View**, and **Table View**. These views allow users to explore data in different formats, depending on their needs.

3.1 View types

Graph View:

- In this view, indicators are displayed as line graphs.
- Users can select indicators and regions using drop-down menus to customize the graph.
- The graph shows predictions for each indicator, including uncertainty levels. A legend is provided to hide or show specific graph lines for easier comparison.
- The vertical axis shows the indicator values, and the horizontal axis represents the time range. Users can zoom into specific time periods by selecting a portion of the graph.

Funded by the European Union

- Each graph includes links to data sources, definitions of the indicator, explanations of the predictions, and additional relevant information.

Map View:

- This view displays indicators on a regional map.
- Users can zoom, pan, and adjust map layers to focus on specific areas.
- Definitions and explanations for the indicators are also available in this view.

Table View:

- This view presents the indicator data in a table format, making it easier to review and compare multiple data points.
- Filters applied in one view (e.g., Graph View) automatically carry over when switching to another view. Tabs at the top of the dashboard make switching between views simple.

Currently, the first prototype uses mock data, as real data is not yet available.

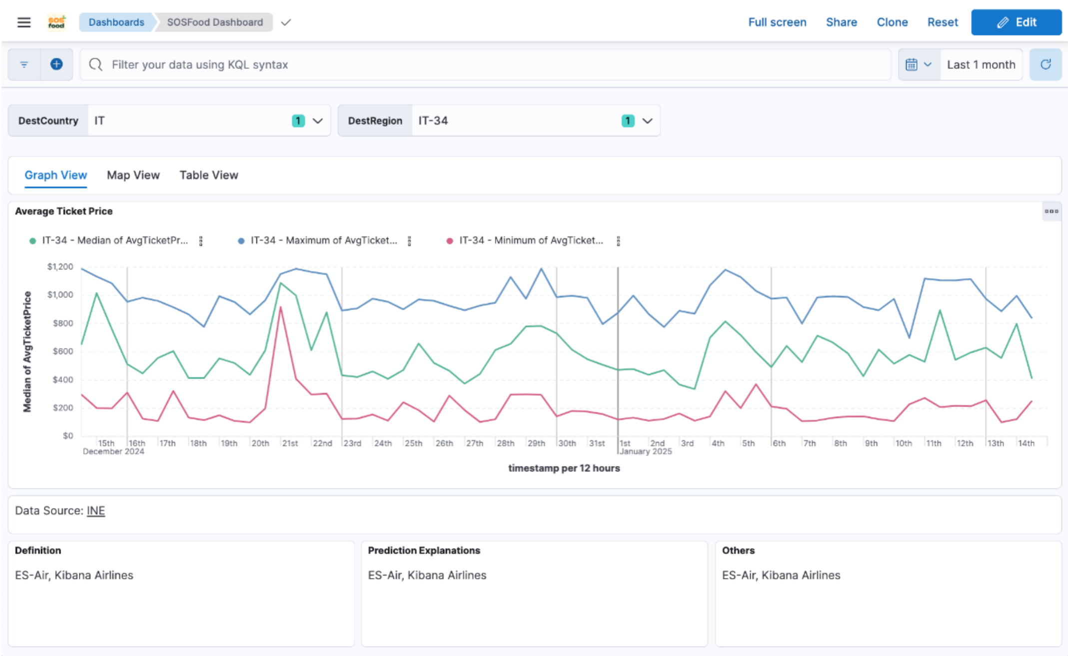

View of the first prototype of the Predictive Dashboard

3.1 Infrastructure

The predictive dashboard is built on a Kibana instance, which has been deployed in a production environment using **Helm charts** and **Kubernetes**. Here's an overview of the setup:

- **Elasticsearch Instance**: An Elasticsearch instance, required by Kibana, was first deployed and configured.
- **AWS and DNS**: The deployment included setting up AWS node instances, DNS configurations, and HTTPS access for security.
- **Kubernetes Configuration**: A Kubernetes service was configured to provide public access, with Helm charts managing the deployment for Elasticsearch and Kibana.
- **User Access**: Kibana was set up to allow both anonymous access for Guest Users and basic login access for users with a username and password.

This infrastructure ensures the dashboard is accessible, secure, and ready for future updates when real data becomes available.

Dashboard Access
URL: https://dashboard.sosfood.jibe.cloud/

4. Testing and validation

The testing and validation phase is essential for the development of the predictive dashboard, to ensure optimal performance before its widespread adoption. The testing is done in controlled and real-world settings through collaboration across work packages and engagement with stakeholders. The goal is to verify its applicability by demonstrating its integration into diverse operational environments and serves several purposes:

- rigorous validation builds confidence in the tool's accuracy and relevance, fostering adoption

- continuous feedback loops ensure the tool evolves to meet user needs and accommodate emerging data or methodologies.

- a validated dashboard presents a tangible outcome that showcases the project's success in driving sustainability through technology.

4.1 The objectives

The testing and validation phase serves several critical objectives:

Funded by the European Union

- **Ensuring Usability**:
  The dashboard must meet high standards of user experience across multiple platforms, including mobile and web.

- **Validating Functionality**:
  Each feature, from data filtering to user roles, needs to perform seamlessly under various scenarios.

- **Guaranteeing Accuracy**:
  The dashboard's visualizations, predictions, and data representations must align with the project's scientific and methodological frameworks.

- **Demonstrating Value**:
  This phase showcases how the dashboard will enhance decision-making for industries, policymakers, researchers, and the general public.

4.2 Timing of the Activities

The testing and validation process unfolds across several phases, ensuring a comprehensive assessment of the dashboard:

February to April 2025: Controlled environment testing (M13—M15). This phase ensures the first prototype meets baseline functionality and usability standards.

May to December 2025: Case study testing and validation (M16—M24).Real-world validation occurs in three geographically and operationally distinct case studies to assess adaptability and scalability.

January to December 2026: Iterative refinement (M25—M36). Feedback from earlier testing phases is used to enhance the dashboard and prepare for broader adoption.

4.3 Use Cases for Testing

To ensure comprehensive validation, several specific use cases have been outlined:

1. **Multiplatform Accessibility**:
   - Validate that the dashboard is mobile-responsive and functions seamlessly on both web and mobile platforms.
   - Ensure a user-friendly interface with consistent performance across devices.

2. **Guest User Access**:
   - Test guest users' ability to view data visualizations and navigate available indicators.
   - Confirm that guest users cannot edit or delete dashboards or visualizations.

Funded by the European Union

- Restrict access to advanced functionalities (e.g., Kibana backend).

3. **Editor User Functionality**:
    - Verify that editor users (Dashboard Designers) can create, edit, and delete dashboards and visualizations.
    - Ensure editor users cannot access functionalities outside their role (e.g., administrative settings).

4. **Data Granularity**:
    - Validate the ability of users to view predictions at different time granularities (daily, weekly, monthly).
    - Confirm smooth toggling between granular views.

5. **Data Filtering and Views**:
    - Test controls that allow users to filter data by parameters (e.g., indicators) and sources (e.g., regions).
    - Validate smooth switching between graph, map, and table views.

6. **Explanatory Information**:
    - Confirm that all visualizations are accompanied by:
        - Definitions of indicators.
        - Explanations of prediction methodologies.
        - Legends and links to data sources.

7. **Data Integrity and Accuracy**:
    - Ensure predictions and visualizations accurately reflect data from WP1 and insights from WP3.
    - Test links to data sources to verify accessibility and validity.

4.4 Controlled environment

The controlled environment testing phase, scheduled from February to April 2025 **(M13–M15)**, ensures that the predictive dashboard's core functionality, usability, and accuracy are validated in a simulated setting before deployment in real-world contexts. Led by SOFTEAM under WP4, this phase involves testing multiplatform accessibility, data integration from WP1, and predictive outputs from WP3 to ensure alignment with project methodologies. Simulated datasets and predefined test cases are used to evaluate features like data filtering, visualization accuracy, and user role restrictions. Feedback from internal stakeholders, including project partners, is collected to address any bugs or usability challenges, creating a refined version of the dashboard ready for real-world validation.

4.5 Real-world environment

The real-world environment testing phase, running from May to December 2025 **(M16–M24)**, validates the dashboard's adaptability, scalability, and relevance in diverse operational contexts across three case studies: Galicia, Athens, and Lithuania. Led by WP5 with contributions from WP1 and WP3, this phase tests the dashboard's integration with live data and its ability to provide actionable insights for stakeholders like policymakers, industries, and researchers. Regional workshops train users to evaluate real-world scenarios, such as sustainability interventions or food system optimizations, and gather feedback on usability and functionality. Insights from these case studies refine the dashboard to ensure it meets regional needs, supports effective decision-making, and is ready for broader adoption

5. Next steps and future roadmap

The development of the Predictive Dashboard is an ongoing process, with several key steps planned to enhance its functionality, integrate additional services, and prepare it for its final deployment. These steps focus on ensuring seamless data ingestion, processing, and visualization to meet the project's goals.

5.1 Data Integration and Ingestion Flow

Currently, the dashboard uses **Kibana** for data visualization, which relies on **Elasticsearch** (from the ELK stack) as its underlying database. To enable the visualization of predictions generated by the AI models developed in WP3, we need to establish a robust data ingestion pipeline. This process involves:

1. **Pushing AI Results to Elasticsearch**: The outputs from WP3's AI services will be integrated into the Elasticsearch database, enabling the dashboard to display predictive insights.

2. **Collecting and Storing Raw Data**: The raw data, which serves as input for the AI models, needs to be collected from various sources and stored in a suitable repository.

3. **Connecting Data Sources to AI Services**: The stored raw data will be sent to WP3's AI services for processing, generating predictions that can be visualized in the dashboard.

Additionally, we may need to develop an **auxiliary service** to handle data collection, ensuring that all relevant datasets are properly ingested and stored. This auxiliary service will support the overall data pipeline and ensure compatibility with the dashboard's requirements.

Funded by the European Union

This proposed data ingestion flow ensures that the Predictive Dashboard remains tightly integrated with the project's broader data and AI ecosystem. However, the final implementation may evolve depending on the progress and outputs of other work packages and any new requirements that emerge for the dashboard. The figure below illustrates this data flow.

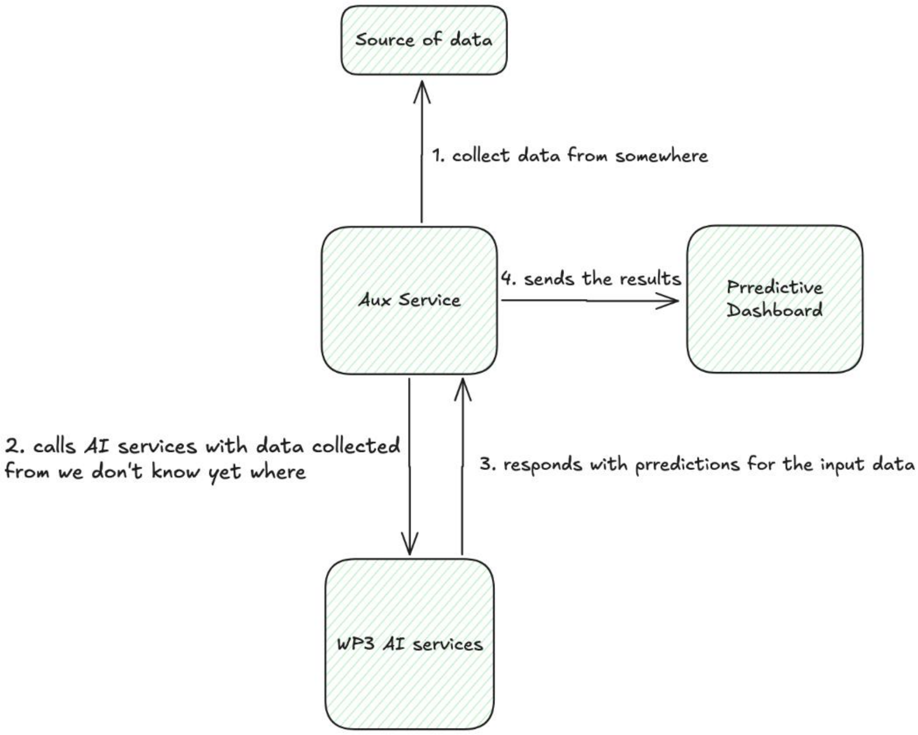

5.2 Future Development and Improvements

In addition to completing the data integration, the following improvements are planned for the Predictive Dashboard:

- **Enhanced Features and Functionalities**:

Based on stakeholder feedback during testing, new features may be added to improve usability and meet specific user needs.

- **Scalability and Performance Optimization**:

As more data becomes available, the dashboard infrastructure will be optimized to ensure it can handle larger datasets efficiently.

Funded by the European Union

- **Customization for Case Studies**:

Specific requirements from the case study regions (Galicia, Athens, and Lithuania) will be incorporated into the dashboard's design and functionality.

- **User Experience Enhancements**:

Iterative refinements will focus on improving navigation, accessibility, and overall user satisfaction.

5.3 Vision for the Final Version

The final version of the Predictive Dashboard will be a fully integrated tool that allows stakeholders to explore data, analyze trends, and make informed decisions based on AI-powered predictions. It will act as a bridge between raw data, advanced analytics, and actionable insights, empowering users across industries, research, and policymaking to drive sustainability in food systems.

Funded by the
European Union