ALLIANCE

A hoListic framework in the quality Labelled food supply chain systems' management towards enhanced data Integrity and verAcity, interoperability, traNsparenCy, and tracEability

Deliverable 2.4
Final AI-enabled tools for vulnerability risk assessment, early warning indication and decision support preventive actions

GRANT AGREEMENT NUMBER: 101084188

This project has received funding from the European Union's HE research and innovation programme under grant agreement No 101084188

Lead Beneficiary: University of Thessaly

Type of Deliverable: Report

Dissemination Level: Public/Confidential

Submission Date: 13.05.2025

Version: 1.0

Versioning and contribution history

Version	Description	Contributions
0.0	Table of Contents	UTH
1.0	Contributions to different sections	UTH
2.0	INTRA, FINS contributions	INTRA, FINS
2.1	Integration of internal reviews	MENA, UNIBO
2.2	UTH final review and integration of modifications	UTH

Authors

Author	Partner
Kostas Choumas	UTH
Apostolis Apostolaras	UTH
Dimitris Kanavaris	UTH
Marialena Lefkopoulou	UTH

Reviewers

Name	Organisation
Tamara Zivadinovic	MENA
Giulia Maesano	UNIBO

Disclaimer

Table of content

List of figures

List of Abbreviations

Abbreviation	Description
AI	Artificial Intelligence
CBV	Core Business Vocabulary
CCP	Critical Control Point
DoA	Description of Action
EPCIS	Electronic Product Code Information Services
EWDSS	Early Warning Decision Support System
FSC	Food Supply Chain
HSI	Hyperspectral Imaging
IoT	Internet of Things
NIR	Near-Infrared
PDO	Protected Designation of Origin
PGI	Protected Geographical Indication
VRAMF	Vulnerability Risk Assessment Management Framework

Executive Summary

This deliverable, D2.4 - “Final AI-enabled tools for Vulnerability Risk Assessment, Early Warning Indication and Decision Support Preventive Actions” provides an updated and comprehensive overview of the implementation status and results achieved by M30 for the relevant technological components developed under WP2. It includes significant updates and advancements reflecting the progress made since the initial submission of the deliverable D2.3. The deliverable details the progress made in the three of the five tasks (T2.3-T2.5) within the work package, focusing on the implementation aspects, and explains the main components and functions that have been successfully implemented so far.

- Task 2.3 - Vulnerability Risk Assessment for Critical Control Points Identification in quality-labelled FSCs
- Task 2.4 – AI-enabled Early Warning and Decision Support System
- Task 2.5 – Interoperability Mechanisms in Complex Food Systems

This final version highlights the following advancements:

- Successful completion of the Early Warning System application, marking the transition to an operationally mature state and showcasing that it is ready for testing and use in real-world scenarios. The key functions implemented have been tested in emulated scenarios, increasing their usefulness for practitioners.
- The Vulnerability Risk Assessment, which reports its evolution to a fully functional and mature stage. It is used to accurately assess and prioritize risks in different operational environments and supports informed decision making.
- The Interoperability System which is based on GS1 standards and EPCIS protocols. The system is now operational and has proven its ability to enable seamless data exchange and interoperability across heterogeneous supply chains, creating transparency, efficiency, and coordination between different actors.
- A reference to the Blockchain Apps and the Knowledge Database as these which are described with a detailed analysis in deliverables D2.2 and D3.2 respectively.

In summary, this deliverable D2.4 provides a comprehensive overview of the results achieved in WP2, focusing on the development and finalisation of the Early Warning and Decision Support System, the Vulnerability Risk Assessment component, and the Interoperability System based on GS1 and EPCIS standards. The results show that these components are ready for integration and deployment in real food supply chain environments. Overall, these systems represent a major step forward in providing smart, interoperable and risk-aware solutions that improve traceability, increase responsiveness to emerging threats and strengthen the resilience and transparency of food quality labelling value chains - in line with the objectives of the project.

1. Introduction

1.1. Document purpose & scope

WP2 is one of the technical core packages of ALLIANCE. It provides key technical components and solutions for the implementation of the ALLIANCE platform. This deliverable D2.4 *Final AI-Enabled Tools for Vulnerability Risk Assessment, Early Warning Indication and Decision Support Preventive Actions* describes the results of WP2 achieved in the period M19-M30 of the project to achieve the following objectives (according to the Description of Action (DoA)):

- **WP2.Obj.1:** To create the Blockchain framework for providing increased traceability in organic, PDO, PGI and GI food products.
- **WP2.Obj.2:** To provide food actors with increased visibility and situational awareness about the performance of the quality labelled Food Supply Chain (FSC) against the strict organic, PDO, PGI and GI standards.
- **WP2.Obj.3:** To design and implement an interoperability framework for consolidating different data sources and incorporating IoT devices and rapid authenticity testing devices.
- **WP2.Obj.4:** To design and implement a Vulnerability Risk Assessment Framework to assess Critical Point within the FSC.
- **WP2.Obj.5:** To design and implement an Early Warning & Decision Support System based on AI and predictive analytics for supporting proactively interventions against food fraud.

The achievement of the above objectives has so far been accomplished through progress in the activities under the following tasks:

- **Task 2.2** - Resilient Food Supply Chain Systems using Blockchain
- **Task 2.3** - Vulnerability Risk Assessment for Critical Control Points Identification in quality-labelled FSCs
- **Task 2.4** – AI-enabled Early Warning and Decision Support System
- **Task 2.5** – Interoperability Mechanisms in Complex Food Systems

The following Table summarizes how each Task has contributed to the WP2 objectives.

Tasks	*Contribution to attain to the WP2 Objectives*
T2.2	**WP2.Obj.1:** Apart from the digital transformation of the current FSCs, the use of the Blockchain Technology in the ALLIANCE platform offers also increased traceability allowing stakeholders to trace back the origin of the food products, verify and justify the data accompany the food products, confirm food sources and ensure quality standards for PDO, PGI, and GI food products. **WP2.Obj.2:** Utilizing Blockchain technology provides transparent and immutable data records, allowing food actors (and according to their roles) access to authenticated and trustworthy information considering the journey of the food products in real-time. Data integrity is ensured through cryptography, which allows information related to production dates, packaging numbers etc. to be accessed only by the authorised users in a secure way.
T2.3	**WP2.Obj.4:** Strengthening the quality labelled FSCs against adulteration and fraud requires continuous and systematic assessment of the related vulnerabilities, the identification of the threats in every step across the entire food chain and the linkage with the associated vulnerabilities, as well as fraud risk assessment. Task 2.4 provides the framework that utilizes the Blockchain technology and the AI with an aim to identify the key point within the FSCs that are prone to vulnerabilities.

T2.4	**WP2.Obj.5:** Leveraging AI technology for analysing the data stemming from the various steps of the FSCs allows for the detection and the identification of behaviours or performance that may indicate a deviation or an anomaly comparing to the normal operation of the FSC, and might be categorized after assessment as a risk or threat which in sequence allows the FSCs operators to take immediate actions to investigate further the incidence and make informed decisions to address it. The ALLIANE AI Early Warning system is fed with real-time data and along with the historic records of the data already being in stored in the Blockchain it is trained to identify unusual patterns or unexpected changes from outliers.
T2.5	**WP2.Obj.3:** As the digital transformation of the food systems and the utilization of different types of Blockchain technology led to the creation of data ecosystems, which are described by heterogeneity in data management and formats, the challenge of data sharing and exchange can lead to fragmented digital chains. By aligning with the GS1 EPCIS standard and adopting a common vocabulary, ALLLIANCE aims to support interoperability among different food systems and facilitate data discovery, sharing and exchange among different food supply chains

1.2. Relation to project work

This deliverable apart from the WP2 objectives is aligned also with the holistic objectives and scope of the project, as it directly contributes to their successful implementation. The documented results reflect the progress toward the specific objectives and demonstrate the technological advancements deemed necessary to support the pilot activities.

Obj.1 To provide food producers and retailers with a holistic framework consisting of innovative methods, state-of-the-art technologies, reliable processes, and interoperable systems that ensure data veracity and accelerate transparency and trust throughout the EU quality-labelled food chains

Obj.2 To investigate the Food Fraud Landscape and propose systemic solutions that move beyond current practices with an aim to enhance traceability, ensure authenticity, preserve quality and eliminate the fraud in food products through novel cost-effective methods and tools that can detect adulteration on the spot and provide trusted interoperable quality-labelled FSCs.

Obj.4 Increase transparency in quality labeled supply chains, of organic, PDO, PGI and GI food, through innovative and improved track-and-trace mechanisms containing accurate, time-relevant, and untampered information for the food product throughout its whole journey from farm to fork.

Obj.5 Equip food actors, farmers, public authorities, and policy makers with meaningful insights to have the complete situational awareness of the food supply chain (in particular organic, PDO, PGI and GI) while at the same time monitoring the financial, nutritional, environmental, social performance of

The final version of the AI-enabled tools for vulnerability risk assessment, early warning indication, and decision support preventive actions is presented in the current deliverable. The document provides comprehensive descriptions of the design of the integrated solutions that have been developed and implemented under Work Package 2 (WP2), specifically through Tasks T2.2 to T2.5. Particularly,

Task T2.2 focused on the development of resilient Food Supply Chain (FSC) systems using Blockchain, enabling secure, transparent, and tamper-proof traceability across quality-labelled agri-food supply chains.

Task T2.3 delivered an advanced Vulnerability Risk Assessment (VRAMF) framework, identifying critical control points in FSCs and risk indicators tailored to specific food products

Task T2.4, an AI-enabled Early Warning and Decision Support System (EWDSS) was implemented, offering predictive capabilities and real-time insights for proactive fraud detection and rapid response across supply chain actors.

Task T2.5 established robust interoperability mechanisms to ensure seamless integration and communication across heterogeneous digital systems and tools within complex food ecosystems, enhancing scalability and data harmonisation.

1.3. Document Structure

The document is structured in 7 major Sections.

Executive summary provides a summary of the whole document.

- **Section 1** introduces the main purpose and scope, the relation to project work and the structure of this deliverable.
- **Section 2** "System Architecture" provides an overview of the ALLIANCE concept and introduces the ALLIANCE Reference Architecture that provides a comprehensive overview encompassing all the different technology solutions of WP2 and WP3.
- **Section 3** "The Blockchain App" provides a brief overview of the ALLIANCE blockchain applications, detailed in D2.2.
- **Section 4** "The Knowledge Database App" provides also a brief overview of the knowledge database applications, detailed in D3.3.
- **Section 5** "The Early Warning and Decision Support System" provides a comprehensive overview of the implementation of the early warning system for timely identification of frauds.
- **Section 6** describes the Vulnerability Risk Assessment for the identification of the critical control points of the food value chains.
- **Section 7** describes the overview to enable interoperability between FSCs.
- Lastly, **Section 8** concludes the document.

2. System Architecture

2.1. Overview

The ALLIANCE architecture consolidates key technologies and data processing layers, such as the **Data Acquisition**, **Data Management** and **Application** layers, as depicted in Figure 1. It is a wholistic approach for FSCs that encompasses the entire process of gathering and utilizing data related to them, from data harvesting to data consumption.

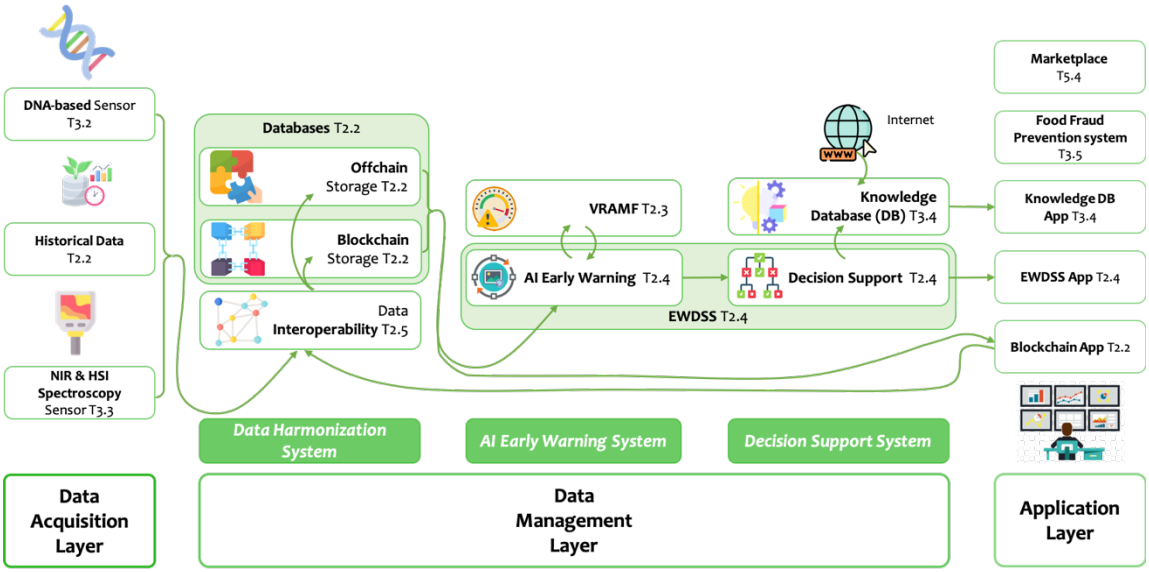

Figure 1: The ALLIANCE Logical Architecture from DoA.

Below, we present the components existing at the three layers of the ALLIANCE architecture, as well as their interactions. All components are mature and completed, apart from the result of T5.4 that will go until M36. We follow a bottom-up approach, according to which:

1. The first layer is the **Data Acquisition** layer. It includes the data sources, which are of three types. It is modular and allows for dynamic extension with additional data sources during the project lifetime or even after its expiration. The two types of data sources are the **DNA-based** and the **NIR & HSI** (Near-Infrared & Hyperspectral Imaging) **Spectroscopy sensors** (results of T3.2 & T3.3, presented in D3.3), and the third type is the **Historical data** that is retrieved from the local databases of the actors involved in the FSCs (result of T2.2).
2. The second layer is the **Data Management** layer, which is responsible for the data processing and consists of three systems: **Data Harmonization, AI Early Warning** and **Decision Support** systems. In turn,
   2.1. The **Data Harmonization** system consists of the **Data Interoperability** process (result of T2.5) that harmonizes the data, which are stored right after in the **Blockchain** and **Off-chain** databases (results of T2.2).
   2.2. The **AI Early Warning** system mainly consists of the **AI Early Warning** process that is the first half of **EWDSS** (Early Warning Decision Support System, result of T2.4). This process is facilitated by **VRAMF** (Vulnerability Risk Assessment Management Framework, result of T2.3). The AI Early Warning process uses the stored data in the Blockchain and Off-chain databases to detect potential food frauds and interacts with **VRAMF**, which continuously exploits the produced warnings to identify the critical control points in the FSCs.

2.3. The **Decision Support** system (result of T2.4) consists of a process of the same name that is the other half of **EWDSS**, which is fed by the AI Early Warning system and suggests actions to the administrator to mitigate the possibilities of food frauds. This system also includes the **Knowledge Database** (result of T3.4), which uses data retrieved by the Blockchain and Off-chain databases and the Internet open datastores to create a broader collection of information that is related to food fraud.

3. Finally, the third layer is the **Application** Layer that includes Mobile/Web Applications, which enable end users to interact with the FSCs. These applications include the **Blockchain App** (result of T2.2), the **Decision Support App** (result of T2.4), the **Knowledge Database App** (result of T3.4), the **Food Fraud Prevention system** (result of T3.5) and the **Marketplace** (result of T5.4). The Blockchain App is used for interacting with the databases, the Decision Support App exports the results of the data analysis, the Knowledge Database App interacts with external sources from the Internet and the Marketplace handles the industrial data.

2.2. Architecture Layers and Components

This section provides a more comprehensive explanation of the three levels of the ALLIANCE architecture and their components. It continues with a more detailed presentation of the ALLIANCE components, providing also references to the following sections for furthermore specific information. Whenever it is necessary, the FSC of Feta Cheese is used as an illustrative example to demonstrate the role of each component.

2.2.1. Data Acquisition Layer

In the Data Acquisition Layer, data is primarily generated and collected automatically through the utilization of distributed IoT sensing devices, rather than being manually injected by users. The generated data either refers to performance metrics from the FCS operations or testing scores of the authenticity and the origin of the food products. Apart from these data collected currently by **DNA-based** and **NIR & HSI Spectroscopy sensors**, there are also **Historical data**, which are datasets of historic metrics from the FSC operations, which are necessary for the data analytics. The Historical data will be updated during the project's lifetime with the information produced by the developed FSCs. The architecture is designed to be flexible and modular, allowing it to easily adapt to any type of IoT device. D3.3 presents in detail the two types of IoT devices that currently are integrated in our architecture. Synthetic datasets have also been generated in the context of Food Prevention System with Predictive Analytics for the Feta Cheese use case.

2.2.2. Data Management Layer

The Data Management Layer is tasked with the storage and processing of data received from the lower layer. It is composed of a centralized service that has the capabilities to store the entire dataset. Additionally, it utilizes a Blockchain distributed ledger for the most critical data. The data are firstly harmonized and then stored in a standardized manner, mitigating their variability and heterogeneity. Simultaneously, there exists a procedure at the same level for utilizing this data to uncover, via AI, methods to improve the performance of the FSCs. The Data Management layer comprises three distinct systems:

a) The **Data Harmonization System**

This system is responsible to harmonize the heterogeneous data coming from different FSCs, allowing their common processing to simplify and enrich their analysis. The data are stored and shared according to the **EPCIS** (Electronic Product Code Information Services) standard of GS1 [1], which is a flagship data sharing standard for enabling visibility within the stakeholders even of different FSCs. EPCIS helps provide the "what, when, where, why and how" of food products, enabling the capture and sharing of interoperable information about their status, location, movement and chain of custody. Together with the **CBV** (Core Business Vocabulary) [2] that is a companion standard to EPCIS, both standards provide definitions of data values that can be used within the data structures used in the data storage.

Part of the data is stored in parallel in the **Blockchain** distributed ledger [3] by leveraging a private permissioned Blockchain network that supports multiple channels, one for each FSC, which can be bridged through cross-chain and data sharing to support interoperability between different FSCs. More details for the utilization of the Blockchain technology are presented later in Section 3. At this point, we would like to highlight that the storage of the whole dataset on Blockchain would be inefficient, since there are big data that could introduce high delays for their Blockchain storage without being critical to be misused or intentionally manipulated. Thus, Blockchain is exclusively used for the storage of the data that needs to be secured, and the centralized storage, called **Off-chain** (as the opposite of Blockchain that is the On-chain database), is used in parallel for the storage of the whole dataset [4].

b) The **AI Early Warning System**

The main component of this system is the **AI Early Warning** process, which is one of the two components of **EWDSS**, the product of T2.4. This process uses AI and the harmonized data to predict and determine with increased probability possible food fraud incidences within the FSCs. It reactively monitors the FSC operational performance to assess the fraud risk factors and the actual fraud vulnerability of the food products. By harnessing the capabilities of AI [5], it proactively recommends interventions, enabling faster and adaptable decision-making processes crucial for mitigating food fraud. As part of the proposed solution, employing a Mamdani Fuzzy Inference System for early warning demonstrates the effectiveness of AI technologies in detecting anomalies within the complex food supply chain. Crucially, this process will be demonstrated in real-life case studies through rigorous testing, with a focus on a practical use case centered around the FSC of Feta Cheese, Organic Honey and Organic Pasta.

VRAMF is a concurrent parallel component that functions as a supplement to the previous process. The result of T3.1, which ended in M6, was the basis for identifying a first set of critical control points [6] in each FSC for mitigating the food fraud incidences. Specifically, each FSC's stakeholders responded to questionnaires, refined through the Delphi technique [7, 8] as it was presented in D2.1, to identify the initial set of critical control points. These control points are mainly the points in each FSC where samples are generated and used for quality control. During the lifetime of ALLIANCE, the effectiveness of the results of the AI process, which relies on the samples and the data produced by the current control points, will be improved by redefining this set of control points. In turn, the change in the control points will affect the AI process, thus, an interacting relationship exists between these two processes. More details on this tool are given in Section 6.

c) The **Decision Support System**

Early warning signals generated by the AI-enabled Early Warning System can serve as critical criteria in the decision-making process, supporting recommendations under

conditions of uncertainty or risk. The Decision Support System plays a dual role: it (i) supports a human-in-the-loop approach where expert opinions and preferences are evaluated for consistency and (ii) facilitates a more automated approach for ranking machine learning algorithms based on performance metrics such as accuracy, precision, and recall. These algorithms are used within ALLIANCE for detecting fraud incidents in the food supply chain. The Decision Support System will be demonstrated through real-world use cases involving various FSCs, including Feta Cheese, Fava Beans, Organic Honey and Organic Pasta, among others. Notably, the human-in-the-loop methodology offers a generic tool applicable across all pilot scenarios.

The **Knowledge Database** is conceptualized as an all-inclusive repository, well-designed with the assimilation of processed data, insights, and inferences derived from the analysis of food products along with their supply chains in an immaculate manner. The integration of external data (standards, certificates, PDO/PGI CoPs, scientific articles, links to related websites, etc.) with the data originating from the project makes it easy to take a thorough examination and extraction of valuable insights and reports by each product. More details are given D3.3.

2.2.3. Application Layer

The Application Layer provides interactive **Web Apps** for comparing and filtering the data analytics and the suggested decisions of the Data Management layer. These user-friendly applications can support multiple roles of end users (such as farmers, producers, processors and retailers), who are informed about the analytics or the decisions of their interest. Moreover, the policy makers and authorities can access this information to design countermeasures for food fraud mitigation.

There are the 4 Apps developed in ALLIANCE:

1. the **Blockchain App**, presented in Section 3 and more detailed in D2.2,
2. the **Knowledge Database App**, presented in Section 4 and more detailed in D3.3,
3. the **EWDSS App**, presented in Section 5 and
4. the **Marketplace**, which will be presented in deliverable D5.6.

3. The Blockchain App

While the technical background of the Blockchain web app is explained in deliverable D2.2, this section provides a summary of the functionality and purpose of the app from both a user and system perspective. The Blockchain web app is used by administrators and staff in the FSCs to monitor supply chains, check the status of food processing, and enter data on their actions. Employees have limited access to the relevant data depending on their role within the FSC. Ultimately, the online application makes it easier to trace the product and present its journey to potential consumers, to reassure them of the authenticity and high quality of the food products. The seven FSCs of ALLIANCE are as follows:

PDO/PGI Extra Virgin Olive Oil (referred from now on as ***Olive Oil*** for simplicity reasons)

PDO Feta Cheese (referred from now on as ***Feta Cheese*** for simplicity reasons)

Organic Honey

PGI Asturian Faba Beans (referred from now on as ***Faba Beans*** for simplicity reasons)

PGI Lika Potatoes (referred from now on as ***Lika Potatoes*** for simplicity reasons)

Organic Pasta

PDO Arilje Raspberries (referred from now on as ***Arilje Raspberries*** for simplicity reasons)

The Blockchain web app increases the resilience of FSCs against various unintended threats or food frauds, which is one of the main goals of ALLIANCE.

The **Blockchain technology is the main pillar of building resilient FCSs**.

Blockchain helps supply chain partners exchange trusted data through approved blockchain solutions. Businesses and consumers want brands to guarantee the authenticity of products, while supply chain participants demand responsible sourcing and better transparency to minimize disputes. Blockchain for FSCs helps supply chain leaders use data to manage disruption and build resilience. Through distributed ledger technology, which provides a shared, single version of the truth, blockchain applications give authorized participants greater insight and transparency into all FSC activities.

Blockchain is a technology that enables transactions, authentications and interactions to be recorded on a network rather than by a single central authority. The innovation of blockchain is that storage does not depend on a central authority collecting all the data, but enables decentralized operations, so that all participants can have their own copy of the stored data. The data is mainly generated by the IoT devices and the human users of the developed apps. Once the data is stored in blockchain, no one can tamper with it.

4. The Knowledge Database App

As part of the infrastructure developed in this deliverable, the Knowledge Database App has been implemented to turn earlier conceptual work into a working solution. While Deliverable 3.3 provides the technical background and detailed architecture, this section provides a summary of the current functionality and purpose of the app from a user and system perspective.

The Knowledge Database App acts as the main access point to the Digital Knowledge Base for Food Fraud, offering stakeholders a structured and user-friendly interface to explore information on fraud risks in the food supply chain. The application was developed with a FastAPI backend and a React-based frontend and deployed on AWS infrastructure to ensure flexibility, security and performance. Thanks to this configuration, various types of data can be processed and displayed, including analytical results, regulatory documents, scientific publications, and certification documents.

A central component of the app is its smart search engine, with filters and categories that allow users to search by product type, fraud category (e.g. dilution, misrepresentation of origin or substitution) and available detection tools. The information is displayed using dynamically generated solution cards that include metadata, links to supporting materials and visual elements for better understanding.

Key features of the Knowledge Database App include:

- An interactive dashboard that presents content organized by product type, fraud category, and intervention strategy.
- Advanced search and filter options for precise access to food fraud cases and tools.
- Metadata integration that connects internal project data with external references, such as scientific literature or fraud alerts.
- Tool cards summarizing fraud detection methods and linking to associated documents.
- Secure login and upload functionalities, enabling partners to contribute data and tag content for consistent structuring.

Importantly, the app is integrated with the EWDSS App developed during this project. The alerts and warnings are displayed directly in the Knowledge Database App, where authorized users can view current alerts and track their progress. This real-time feedback loop transforms the app from a static knowledge source into an active decision-support tool.

The app's design emphasizes accessibility and supports a wide range of users, including food industry professionals, regulators and researchers. Its modular architecture allows for ongoing updates and enhancements as new use cases or datasets become available.

Currently, the Knowledge Database App is being validated through practical case studies, including the raspberries supply chain, olive oil traceability, and faba bean case. These pilots will help to refine the functionality and demonstrate how the tool supports fraud risk analysis in real-life situations. It is expected that the system will be further expanded in the future to include other products covered by ALLIANCE.

In summary, the Knowledge Database App represents a key output of the platform, offering a robust and practical tool that brings together structured knowledge, detection technologies, and real-time risk alerts to support fraud mitigation efforts across the agri-food sector.

5. The EWDSS App

5.1. Overview

The development of the ALLIANCE AI-enabled EWDSS aimed at enhancing the detection and mitigation of fraud incidents in the food supply chain is a subject of Task 2.4, led by Netcompany-Intrasoft. As detailed in Deliverable D2.3 "Interim AI-enabled tools for Vulnerability Risk Assessment, Early Warning Indication and Decision Support Preventive Actions", submitted in M18, the objective is to combine Artificial Intelligence methods with Operations Research techniques to generate early warning signals based on **Fuzzy Logic** concepts, and to deliver actionable recommendations through multi-criteria decision analysis approaches. In the context of the current Deliverable D2.4, we showcase the application of the early warning and decision support modules of the AI-enabled EWDSS across several use cases, including Feta Cheese, Organic Pasta, Organic Honey, and Fava Beans.

Furthermore, we extend the Decision Support module by introducing a new multi-criteria decision analysis tool selected to evaluate machine learning algorithms used for fraud detection in the food supply chain, based on multiple performance criteria (e.g., Accuracy, F1-score, Precision, Recall, etc.). We assume that several machine learning algorithms are trained to detect fraud in the food supply chain. Our goal is to support the selection of the most suitable model for use in an operational environment. To achieve this, we implement a champion-challenger framework, in which multiple algorithms are compared and simultaneously evaluated against multiple performance metrics. The use cases of Organic Honey and Fava Beans are used as illustrative examples for this evaluation.

5.2. Early Warning System for Food Fraud Detection

5.2.1. Application in the Organic Pasta Supply Chain

As part of the Early Warning System, we explore the use of *fuzzy logic* to assess the risk of fraud within the Organic Pasta supply chain. The aim is to increase the transparency and reliability of organic certification processes by using linguistic variables and aggregating three critical input factors.

First, we perform a multi-residual analysis to assess the level of contamination in three stages of the production process: wheat, semolina and pasta. These levels are assessed against a threshold of 0.01 ppm to ensure compliance. Based on the contamination concentration, each batch is categorized into one of three fuzzy risk levels: Low (indicating minimal contamination and safety for certification), Medium (indicating possible irregularities) and High (flagging unacceptable contamination and a likely fraud scenario). Second, we analyze the field audit compliance percentage, which indicates how well farms comply with organic farming standards. Farms with a compliance rate of less than 52% are automatically flagged and excluded from certification. The level of compliance is divided into the categories: Poor (below 52%), Average (between 52% and 80%), and Good (above 80%). This input is an important indicator of the overall risk situation of the company. The third input involves assessing the validity of organic certifications issued by actors throughout the supply chain, including pasta factories, mills, farmers and agricultural advisors. Each expired certification leads to an automatic rejection of the respective supplier. The certification status is categorized in fuzzy sets: Valid (indicating a low risk of fraud) and Expired (indicating a high risk of fraud). These

inputs are processed by the fuzzy logic engine, which calculates a fraud score within the range [0, 100] and assigns a corresponding fraud label (Low, Medium, or High).

The following rules govern the behavior of the system:

- **Rule 1**: If any of the contamination levels in pasta, semolina, or wheat is *high*, then the fraud risk is considered *high*.
- **Rule 2**: If the field audit compliance is *poor* **and** any of the organic certifications (from factory, miller, farmers, agronomic studio, or brand) is *expired*, then the fraud risk is *high*.
- **Rule 3**: If all contamination levels are *low*, all certifications are *valid*, **and** the audit compliance is *good*, then the fraud risk is *low*.
- **Rule 4**: If any of the certifications is *expired* (regardless of other factors), the fraud risk is *medium*.
- **Rule 5**: If the audit compliance is *average*, the fraud risk is also set to *medium*.

These rules form the basis of the Fuzzy Logic control system that determines the fraud score. Inputs from real-world cases are mapped to fuzzy sets, and the system aggregates these through a Mamdani-style inference engine (for details see Deliverable D2.3). The resulting score is defuzzified into a crisp value and then interpreted as *low*, *medium*, or *high* fraud risk based on thresholding logic: scores below 40 are labeled *low*, scores between 40 and 70 as *medium*, and scores above 70 as *high*.

To enable seamless integration and practical use of the Early Warning System in real-world organic pasta scenarios, we have deployed the fraud detection solution as a RESTful API, accessible via the following URL: FastAPI - Swagger UI. This API allows users to submit fraud assessment requests in JSON format and receive automated fraud risk evaluations in real time.

Each request includes a list of cases where each case provides the following input features:

- Multi-residual analysis values for *pasta*, *semolina*, and *wheat* (in ppm),
- Field audit compliance percentage (ranging from 0 to 100),
- Binary indicators (1 = valid, 0 = expired) for organic certifications from various stakeholders (*pasta factory*, *miller*, *farmers*, *agronomic studio*, and *Alce Nero*).

The API processes these inputs through the fuzzy inference engine and returns a JSON response containing:

- The case ID,
- A fraud score (numeric value between 0 and 100),
- A fraud label ("low", "medium", or "high"),
- A URL pointing to a plot that visualizes how the fraud risk was derived based on fuzzy membership functions.

Tables 1 and 2 present JSON request and JSON response examples.

```
[
 {
  "case_id": "128",
```

```
  "multiresidual_analysis_pasta": 0.003,
  "multiresidual_analysis_semolina": 0.005,
  "multiresidual_analysis_wheat": 0.002,
  "field_audit_percentage": 90,
  "organic_certification_pasta_factory": 1,
  "organic_certification_pasta_miller": 1,
  "organic_certification_farmers": 1,
  "organic_certification_agronomic_studio": 0,
  "organic_certification_alce_nero": 1
 },
 {
  "case_id": "129",
  "multiresidual_analysis_pasta": 0.01,
  "multiresidual_analysis_semolina": 0.008,
  "multiresidual_analysis_wheat": 0.012,
  "field_audit_percentage": 65,
  "organic_certification_pasta_factory": 0,
  "organic_certification_pasta_miller": 1,
  "organic_certification_farmers": 0,
  "organic_certification_agronomic_studio": 1,
  "organic_certification_alce_nero": 0
 }
]
```


Table 1: Example of a JSON request (use case of organic pasta)


```
[
  {
    "case_id": "128",
    "fraud_score": 50.0,
    "fraud_label": "medium",
    "plot_url": "/static/plot_128.png"
  },
  {
    "case_id": "129",
```

```
    "fraud_score": 59.37,
    "fraud_label": "medium",
    "plot_url": "/static/plot_129.png"
  }
]
```


Table 2: Example of a JSON response (use case of organic pasta)

The following Figure 2 provides a clear visualization of the membership functions for both input and output variables, illustrating how the crisp input values (i.e., current values) are mapped to their corresponding fuzzy sets and how the fraud risk is derived as both a fuzzy and a defuzified (i.e., crisp) output value. This rule-based reasoning framework allows for explainable AI behavior, as each decision can be traced back to human-understandable logic.

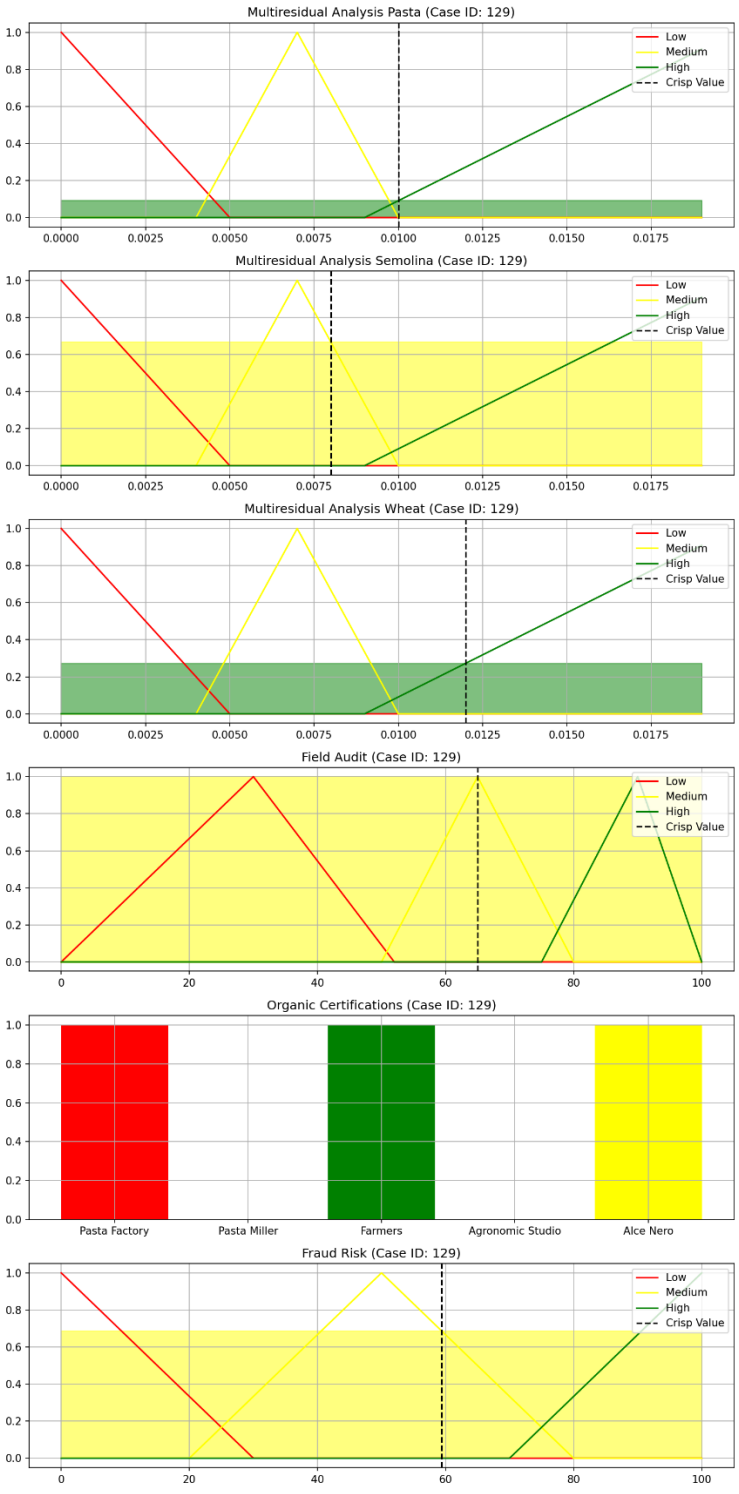

Figure 2: Membership functions of input and output variables in the fuzzy fraud detection model

5.2.2. Application in the Feta Cheese Supply Chain

The early warning system is also used to assess fraud risk levels in the feta cheese supply chain by analysing key parameters such as pH values and temperature conditions. As a Protected Designation of Origin (PDO) product, feta cheese must meet strict quality standards, and any deviation in pH or temperature could indicate potential fraud. The system enables stakeholders to identify suppliers at risk before raw milk, such as goat or sheep milk, is mixed with other milk batches and delivered to the factory for feta cheese production. Therefore, the primary goal is to analyse these critical parameters in real-time, assigning a fraud risk level to each supplier and allowing stakeholders to detect potential fraud or quality issues early in the process.

To support real-time risk monitoring in the feta cheese supply chain, we developed a RESTful API that exposes the early warning system as a service. The API receives structured JSON data as input, evaluates it using fuzzy logic, and returns a risk classification for each milk supplier. The API is accessible via Food Fraud Fuzzy Logic API - Swagger UI (see Figure 3).

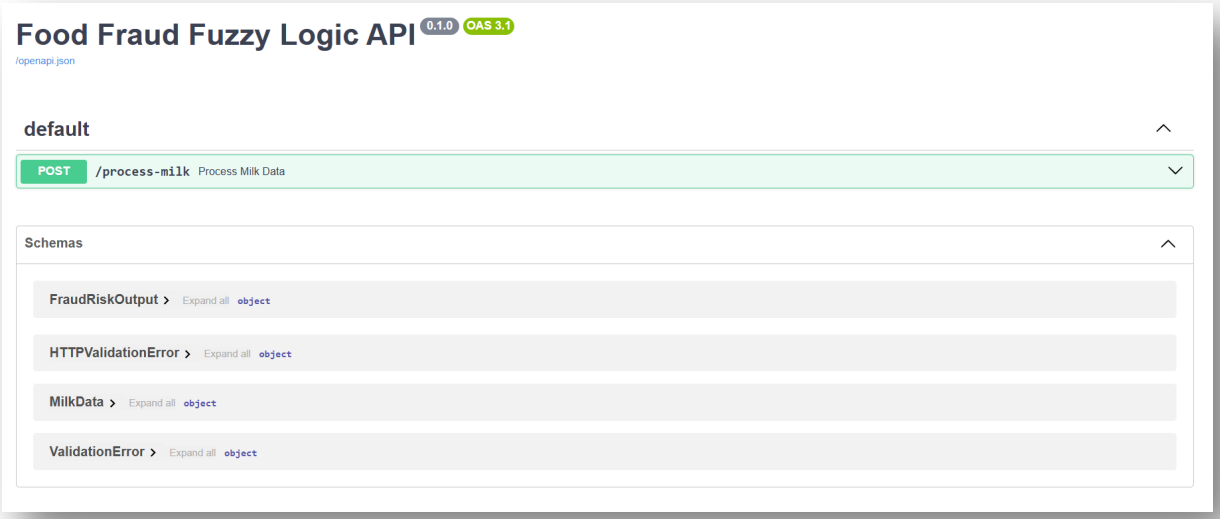

Figure 3: Early warning system as a service

Each request includes records with the following crisp input values (see Table 3):

- supplier_id: Unique identifier for each supplier.
- ph_value: The measured pH of the milk. Feta cheese typically requires pH values between 4.4 and 4.9, but in this context (before production), higher-than-normal values (e.g., > 6.6) may indicate adulteration, contamination, or poor handling.
- temperature_value (°C): The storage temperature of the milk. Ideal values should remain at or below 4°C, as higher temperatures can compromise freshness and suggest cold chain non-compliance.

The system transforms these crisp values into fuzzy sets using triangular and trapezoidal membership functions. In particular, the pH values are categorized as bad, medium, or good, and temperature values are similarly classified. A set of fuzzy rules then evaluates combinations of these conditions to determine a risk level:

- **Rule 1:** If pH is *bad* **or** temperature is *bad*, then risk level is *high (i.e., high risk is inferred when either the acidity is concerned, or the storage temperature is too high).*
- **Rule 2:** If pH is *bad* **or** temperature is *good*, then risk level is *high (i.e., even if the temperature is acceptable, poor pH alone triggers a high risk).*
- **Rule 3:** If pH is *good* **or** temperature is *bad*, then risk level is *medium (i.e. a good pH doesn't fully compensate for a poor temperature, leading to moderate risk).*
- **Rule 4:** If pH is *good* **or** temperature is *medium*, then risk level is *low (i.e., either good acidity or moderately acceptable temperature results in low risk).*
- **Rule 5:** If pH is *medium* **or** temperature is *good*, then risk level is *high (i.e., medium pH still raises concern even with good temperature).*
- **Rule 6:** If pH is *good* **and** temperature is *good*, then risk level is *low (i.e., both quality indicators are within ideal range, confirming low risk).*
- **Rule 7:** If pH is *bad* **and** temperature is *medium*, then risk level is *high (i.e., when acidity is poor and temperature is borderline, the system flags high risk).*

After applying the fuzzy rules, the system produces a continuous risk score in the range of 1 to 3. To simplify interpretation for end-users, a thresholding function is applied to convert the fuzzy output into discrete risk levels. If the risk score is below 1.5, the supplier is classified as *low* risk (level 1); a score between 1.5 and 2.5 corresponds to *medium* risk (level 2); and scores equal to or above 2.5 are labelled as *high* risk (level 3).

The API responds with a JSON-formatted list (see Table 4), where each supplier entry includes the original identifier, the computed risk level (numeric), and its corresponding label ("Low", "Medium" or "High").


```
[
 { "supplier_id": "S001", "ph_value": 6.66, "temperature_value": 5 },
 { "supplier_id": "S005", "ph_value": 6.67, "temperature_value": 4 },
 { "supplier_id": "S006", "ph_value": 6.75, "temperature_value": 2 },
 { "supplier_id": "S008", "ph_value": 6.62, "temperature_value": 8 }
]
```


Table 3: Example of a JSON request (use case of feta cheese)


```
[
  {
    "supplier_id": "S001",
    "risk_level": 2,
    "risk_label": "Medium"
  },
  {
    "supplier_id": "S005",
    "risk_level": 3,
    "risk_label": "High"
  },
```

```
  {
    "supplier_id": "S006",
    "risk_level": 2,
    "risk_label": "Medium"
  },
  {
    "supplier_id": "S008",
    "risk_level": 3,
    "risk_label": "High"
  }
]
```


Table 4: Example of a JSON response (use case of feta cheese)

5.3. Decision Support System for Food Fraud Management

5.3.1. Multi-criteria Evaluation of Machine Learning Algorithms

This section presents a decision support framework based on the TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) method, an established and widely recognized technique in multi-criteria decision analysis [17]. The process begins with the creation of a weighted decision matrix, in which each alternative is evaluated against multiple criteria, each reflecting a different dimension of performance or value. These criteria are normalized and adjusted according to user-defined weights to ensure a fair and meaningful comparison. The method then calculates the geometric distance of each alternative from an ideal solution (defined by the best achievable performance across all criteria) and from an anti-ideal solution (representing the worst possible outcomes). Finally, the distances of each alternative from the ideal and the anti-ideal solutions are compared, allowing the calculation of a similarity coefficient for each alternative. Based on these coefficients, the alternatives are then evaluated and ranked accordingly. The TOPSIS algorithm consists of the following 7 steps:

Step 1: Construct the Decision Matrix

Initially, it is assumed that a decision matrix is established, consisting of m machine learning algorithms, $ML_1, ML_2, \dots, ML_m$, and evaluated based on n performance criteria, $PC_1, PC_2, \dots, PC_n$. Each machine learning algorithm is assessed with respect to each performance criterion individually, forming a decision matrix $X = \left[x_{ij}\right]_{m \times n}$. Additionally, let $W = (w_1, w_2, \dots, w_n)$ denote the vector of criteria weights, where the sum of all weights is equal to 1, i.e., $\sum_{j=1}^{n} w_j = 1$. For each machine learning algorithm i, performance criteria are organized into a decision matrix X, where each entry x_{ij} represents the value of performance criterion j for machine learning algorithm i.

$$X = \begin{bmatrix} x_{11} & x_{12} & \cdots & x_{1n} \\ x_{21} & x_{22} & & x_{2n} \\ \vdots & & \ddots & \vdots \\ x_{m1} & x_{m2} & \cdots & x_{mn} \end{bmatrix} \quad (1)$$

Step 2: Normalize the Decision Matrix

To facilitate the comparison of performance criteria with different rages, the initial step involves normalizing the data. This results in the normalized matrix R:

$$R = \begin{bmatrix} r_{11} & r_{12} & \cdots & r_{1n} \\ r_{21} & r_{22} & & r_{2n} \\ \vdots & & \ddots & \vdots \\ r_{m1} & r_{m2} & \cdots & r_{mn} \end{bmatrix} \quad (2)$$

The normalized value r_{ij} for each performance criterion x_{ij} is computed using the formula:

$$r_{ij} = \frac{x_{ij}}{\sqrt{\sum_{i=1}^{m} x_{ij}^2}}, \qquad i = 1, \dots, m, \qquad j = 1, \dots, n \quad (3)$$

Step 3: Calculation of the Weighted Normalized Decision Matrix

In TOPSIS, the weights assigned to performance criteria are the only subjective parameters, typically reflecting the data scientist's judgment regarding the importance of each criterion in the machine learning lifecycle. The next step is to apply these weights to the normalized decision matrix. The weighted normalized values are calculated by multiplying each normalized value r_{ij} by its corresponding weight w_j:

$$v_{ij} = w_j r_{ij}, \qquad i = 1, \dots, m, \qquad j = 1, \dots, n \quad (4)$$

Step 4: Determine Ideal and Anti-Ideal Solutions

The ideal solution A^+ represents the most favorable outcome for each performance criterion (e.g., high accuracy or low loss), while the anti-ideal solution A^-represents the least favorable outcome (e.g., low accuracy or high loss). The ideal solution A^+ is calculated as:

$$A^+ = (v_1^+, v_2^+, \dots, v_n^+) = \left(\left(\max_j v_{ij} | i \in I_b \right), \left(\min_j v_{ij} | i \in I_c \right) \right), \qquad i = 1, \dots, m, \qquad j = 1, \dots, n \quad (5)$$

Similarly, the anti-ideal solution A^- is determined as:

$$A^- = (v_1^-, v_2^-, \dots, v_n^-) = \left(\left(\min_j v_{ij} | i \in I_b \right), \left(\max_j v_{ij} | i \in I_c \right) \right), \qquad i = 1, \dots, m, \qquad j = 1, \dots, n \quad (6)$$

Here, I_b is associated with criteria considered as benefits, and I_c, with criteria considered as costs.

Step 5: Calculate the Separation Measures

In this step, the distance of each machine learning algorithm from the ideal and anti-ideal solutions is calculated. The Euclidean distance from the ideal solution is:

$$D_i^+ = \sqrt{\sum_{j=1}^{n} \left(v_{ij} - v_j^+ \right)^2} \quad (7)$$

Similarly, the Euclidean distance from the anti-ideal solution is:

$$D_i^- = \sqrt{\sum_{j=1}^{n} \left(v_{ij} - v_j^-\right)^2} \qquad (8)$$

Step 6: Compute the Relative Closeness to the Ideal Solution

The relative closeness C_i^+ to the ideal solution is then computed as follows:

$$C_i^+ = \frac{D_i^-}{D_i^+ + D_i^-}, \qquad 0 \le C_i^+ \le 1, \qquad i = 1, \ldots, m \qquad (9)$$

Step 7: Rank Machine Learning Algorithms

Finally, machine learning algorithms are ranked based on their relative closeness C_i^+ to the ideal solution, with higher values indicating a higher priority for deployment in the operational environment.

Evaluation in the Fava Beans Supply Chain

We apply the multi-criteria decision analysis approach mentioned in Sub-section 4.3.1 to a use case in the fava bean supply chain, focused on detecting fraudulent products using Near-Infrared (NIR) spectroscopy combined with machine learning. The underlying task is a binary classification problem, aiming to distinguish between authentic fava beans from Asturias (representing no fraud) and fraudulent ones from Bolivia (representing fraud).

To enhance signal quality and reduce irrelevant variability in the spectral data, two pre-treatment techniques were applied: Standard Normal Variate (SNV) to correct for scatter effects, and detrending to eliminate linear or quadratic trends in the data. These pre-processing steps were followed by Partial Least Squares Discriminant Analysis (PLS-DA), which was used for dimensionality reduction and feature extraction.

Three classifiers were then applied to the PLS-DA-transformed data: XGBoost (XGB); a Softmax classifier; and a Support Vector Machine (SVM). The study evaluated six modeling pipelines, each combining a specific pre-treatment with a classifier:

- ML1: SNV + PLS-DA + XGB
- ML2: Detrend + PLS-DA + XGB
- ML3: SNV + PLS-DA + Softmax
- ML4: Detrend + PLS-DA + Softmax
- ML5: SNV + PLS-DA + SVM
- ML6: Detrend + PLS-DA + SVM

Each pipeline was assessed using a comprehensive set of performance metrics:

- PC1: Cross-Validation (CV) Score
- PC2: Accuracy
- PC3: Precision (Asturias class)
- PC4: Precision (Bolivia class)
- PC5: Recall (Asturias class)
- PC6: Recall (Bolivia class)
- PC7: F1-score (Asturias class)

- PC8: F1-score (Bolivia class)
- PC9: Area Under the ROC Curve (AUC).

The main idea is to apply the TOPSIS algorithm for evaluating the six modelling pipelines, developed by ASINCAR (ALLIANCE partner), with respect to defined performance metrics. For this purpose, we developed a RESTful API that exposes the TOPSIS algorithm as a service. The API receives structured JSON data as input (see Table 5 for data used), evaluates it using TOPSIS algorithm, and returns a score for each modelling pipeline (see Table 6 for scores obtained). The API is accessible via FastAPI - Swagger UI.

```
{
 "matrix": [
  [0.8565, 0.8657, 0.8571, 0.8763, 0.8947, 0.8333, 0.88, 0.85, 0.864],
  [0.8632, 0.8657, 0.8571, 0.8763, 0.8947, 0.8333, 0.88, 0.85, 0.864],
  [0.8352, 0.8796, 0.8729, 0.8878, 0.8947, 0.8333, 0.89, 0.87, 0.9217],
  [0.8296, 0.8519, 0.8596, 0.8431, 0.8596, 0.8431, 0.86, 0.84, 0.9152],
          [0.834, 0.8611, 0.8684, 0.8529, 0.8684, 0.8529, 0.87, 0.85, 0.9223],
          [0.8374, 0.8611, 0.875, 0.8461, 0.8596, 0.8627, 0.87, 0.85, 0.925]
 ],
 "weights": [0.10, 0.05, 0.10, 0.15, 0.10, 0.20, 0.05, 0.15, 0.10],
 "norm_method": "v",
 "ideal_solution_method": "m",
 "plot_results": false,
 "names_of_alternatives": ["ML1", "ML2", "ML3", "ML4", "ML5", "ML6"],
 "names_of_criteria": ["PC1", "PC2", "PC3", "PC4", "PC5", "PC6", "PC7", "PC8", "PC9"]
}
```

Table 5: JSON request for the TOPSIS service

The JSON request body is structured to include the decision matrix, weights, and other parameters necessary for the TOPSIS methodology. Here's a breakdown:

- matrix: This 2D array represents the decision matrix where each row corresponds to an alternative, and each column corresponds to a criterion.
- weights: A list of values representing the importance of each criterion.
- norm_method: Defines the normalization method to be used. "v" indicates vector normalization.
- ideal_solution_method: Specifies the method to determine the ideal solution. "m" indicates a method based on maximizing or minimizing values to define the ideal.
- plot_results: A boolean value indicating whether visual analytics, such as plots and charts, should be included in the response.
- names_of_alternatives: A list of alternative names.

- names_of_criteria: A list of criterion names.

The JSON response provides the calculated closeness coefficient for each alternative, which indicates how closely each alternative approaches the ideal solution.


```
{
  "closeness_coefficient": [
    {
      "alternative": "ML1",
      "score": 0.4280261695384979
    },
    {
      "alternative": "ML2",
      "score": 0.4280261695384979
    },
    {
      "alternative": "ML3",
      "score": 0.617914617061615
    },
    {
      "alternative": "ML4",
      "score": 0.298935204744339
    },
    {
      "alternative": "ML5",
      "score": 0.44151821732521057
    },
    {
      "alternative": "ML6",
      "score": 0.5060983300209045
    }
  ],
  "execution_time": 6.171,
  "message": "TOPSIS calculation completed successfully."
}
```


Table 6: JSON response for the TOPSIS service

The service does not stop at numerical output (i.e. scores). It includes a suite of visual analytics to enhance interpretability and transparency. Detailed line plots and radar charts illustrate how each alternative performs relative to others, highlighting strengths and weaknesses briefly (see Figure 4 and Figure 5). Bar graphs break down the scores criterion by criterion, offering a deeper understanding of the decision drivers (see Figure 4). Annotated ranking plots and directed graphs then summarize the overall results, guiding stakeholders clearly toward the best-informed decision (see Figure 5 and Figure 6).

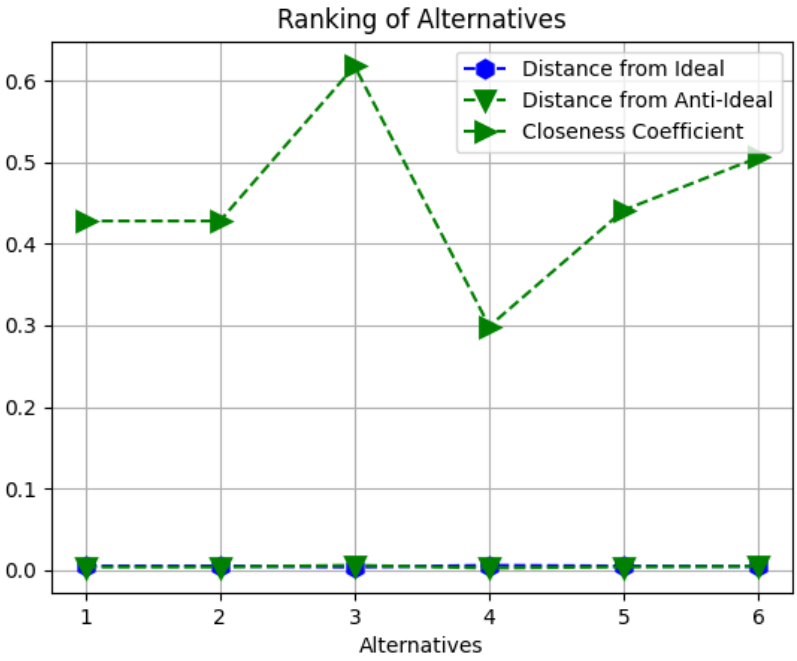

Figure 4: Line plot for illustrating how each alternative performs relative to others

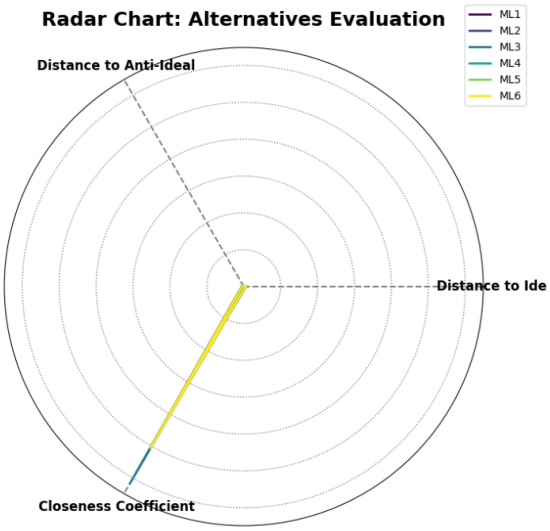

Figure 5: Radar chart illustrating how each alternative performs relative to others

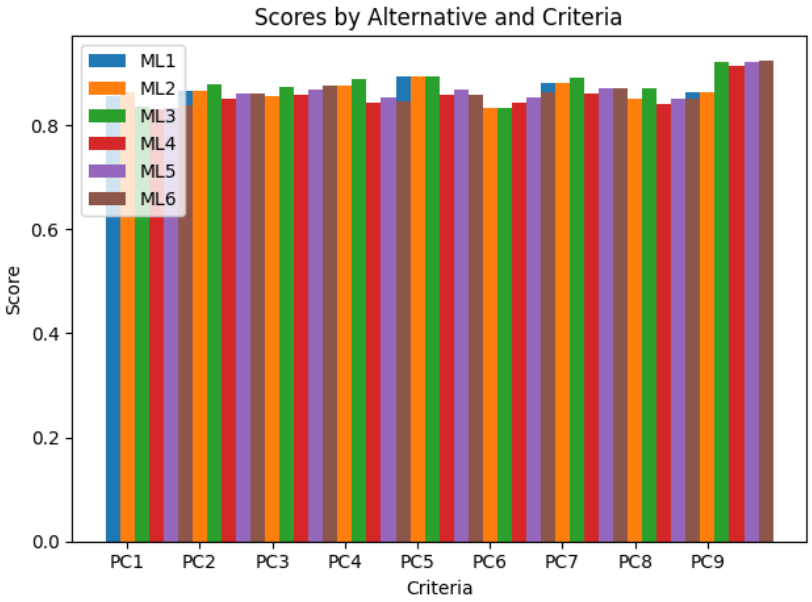

Figure 6: Bar graph for breaking down the scores, criterion by criterion

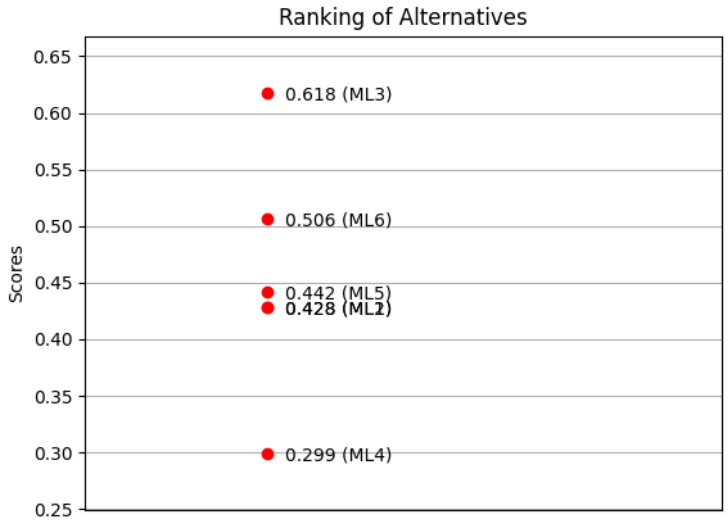

Figure 7: Annotated ranking plot for summarizing the overall results

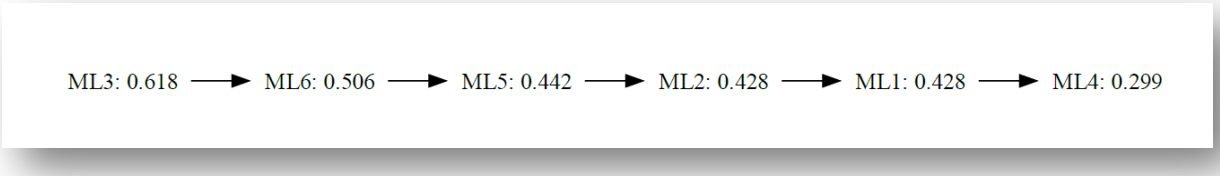

Figure 8: Directed graph for summarizing the overall results

In scenarios where criteria may be conflicting, an alternative version of the TOPSIS service can be used (accessible via: TOPSIS API - Swagger UI). This version allows the client to specify, via the JSON request, whether each criterion represents a benefit (profit) or a cost.

The following JSON request illustrates (see Table 7) how to submit a set of alternatives and criteria to the TOPSIS API. The matrix field contains the performance data for each alternative across the specified criteria. Each row corresponds to an alternative, and each value represents the performance of that alternative on a given criterion. The weights field assigns relative importance to each criterion, where values should sum to 1. In this case, the weights indicate that criterion PC6 (with a weight of 0.20) is the most important, while criteria like PC1 and PC2 have lesser weight. The signs field specifies whether a criterion is a benefit criterion or a cost criterion. A "+" sign indicates that higher scores for that criterion are preferable (e.g., accuracy, AUC), while a "-" sign marks a cost criterion, where lower values are better (e.g., loss, error rates). In this example, all criteria except PC10 are marked as benefit criteria, with PC10 representing a loss value (i.e., a cost criterion), which the model should minimize.

The names_of_alternatives field contains the labels for each alternative (ML1 through ML6). The names_of_criteria field contains the labels for the criteria evaluated before (PC1, ..., PC9) plus a new criterion related to the loss (PC10).

```
{
 "matrix": [
  [0.8565, 0.8657, 0.8571, 0.8763, 0.8947, 0.8333, 0.88, 0.85, 0.864, 4.8392],
  [0.8632, 0.8657, 0.8571, 0.8763, 0.8947, 0.8333, 0.88, 0.85, 0.864, 4.8392],
  [0.8352, 0.8796, 0.8729, 0.8878, 0.8947, 0.8333, 0.89, 0.87, 0.9217, 0.3599],
  [0.8296, 0.8519, 0.8596, 0.8431, 0.8596, 0.8431, 0.86, 0.84, 0.9152, 0.3785],
          [0.834, 0.8611, 0.8684, 0.8529, 0.8684, 0.8529, 0.87, 0.85, 0.9223, 0.3573],
          [0.8374, 0.8611, 0.875, 0.8461, 0.8596, 0.8627, 0.87, 0.85, 0.925, 0.3496]
 ],
 "weights": [0.05, 0.05, 0.10, 0.15, 0.10, 0.20, 0.05, 0.15, 0.10, 0.05],
 "signs": ["+","+","+","+","+","+","+","+","+","+", "-"],
 "names_of_alternatives": ["ML1", "ML2", "ML3", "ML4", "ML5", "ML6"],
 "names_of_criteria": ["PC1", "PC2", "PC3", "PC4", "PC5", "PC6", "PC7", "PC8", "PC9", "PC10"]
}
```

Table 7: JSON request for the TOPSIS service (in scenarios where criteria may be conflicted)

The following JSON output (see Table 8) represents the result of a TOPSIS calculation based on the input provided. The response includes several key fields that contain the results of the scoring, as well as the paths to the generated visual charts.

- closeness_coefficient: This field contains a list of alternatives along with their corresponding TOPSIS closeness scores. These scores indicate how close each alternative is to the ideal solution, where a higher score represents a better alternative.

- execution_time_ms: This field indicates the time it took to perform the TOPSIS calculation, in milliseconds. In this case, the calculation took approximately 637.897 milliseconds.
- message: This field provides a status message about the completion of the TOPSIS calculation. In this example, the message indicates that the calculation was successfully completed.
- bar_chart_path: This field contains the path to the generated bar chart image, which shows the TOPSIS scores for each alternative in a bar graph format. This chart helps users quickly assess the performance of different alternatives in relation to each other. The path is relative to the /charts endpoint, meaning users can access the chart by navigating to this path in their web browser. It can be accessed via the path /charts/bar_chart_e2fa7119ebf248b89981794d13c75481.png (see Figure 9).
- radar_chart_path: This field contains the path to the generated radar chart image that visualizes the performance of alternatives across all criteria. The chart helps users visually compare the alternatives. In this case, the radar chart can be found at /charts/radar_chart_6e7273765b7f4eaa8ef621e9cc684156.png (see Figure 10).
- ranked_line_chart_path: This field contains the path to the generated ranked line chart, which visualizes the ranked alternatives in terms of their TOPSIS scores. The line chart shows the alternatives in descending order of their scores, helping users easily identify the best alternative. The path to this chart is /charts/ranked_line_chart_7083de62df6b4b26963128a7a6e49dd4.png (see Figure 11).


```
{
  "closeness_coefficient": [
    {
      "alternative": "ML1",
      "score": 0.3803081927197713
    },
    {
      "alternative": "ML2",
      "score": 0.38415833297044555
    },
    {
      "alternative": "ML3",
      "score": 0.5790606110459441
    },
    {
      "alternative": "ML4",
      "score": 0.2863103132983457
```

```
    },
    {
      "alternative": "ML5",
      "score": 0.49653788411336164
    },
    {
      "alternative": "ML6",
      "score": 0.5521518530445959
    }
  ],
  "execution_time_ms": 637.897,
  "message": "TOPSIS calculation completed successfully.",
  "radar_chart_path": "/charts/radar_chart_6e7273765b7f4eaa8ef621e9cc684156.png",
  "bar_chart_path": "/charts/bar_chart_e2fa7119ebf248b89981794d13c75481.png",
  "ranked_line_chart_path": "/charts/ranked_line_chart_7083de62df6b4b26963128a7a6e49dd4.png"
}
```


Table 8: JSON response for the TOPSIS service (in scenarios where criteria may be conflicted)

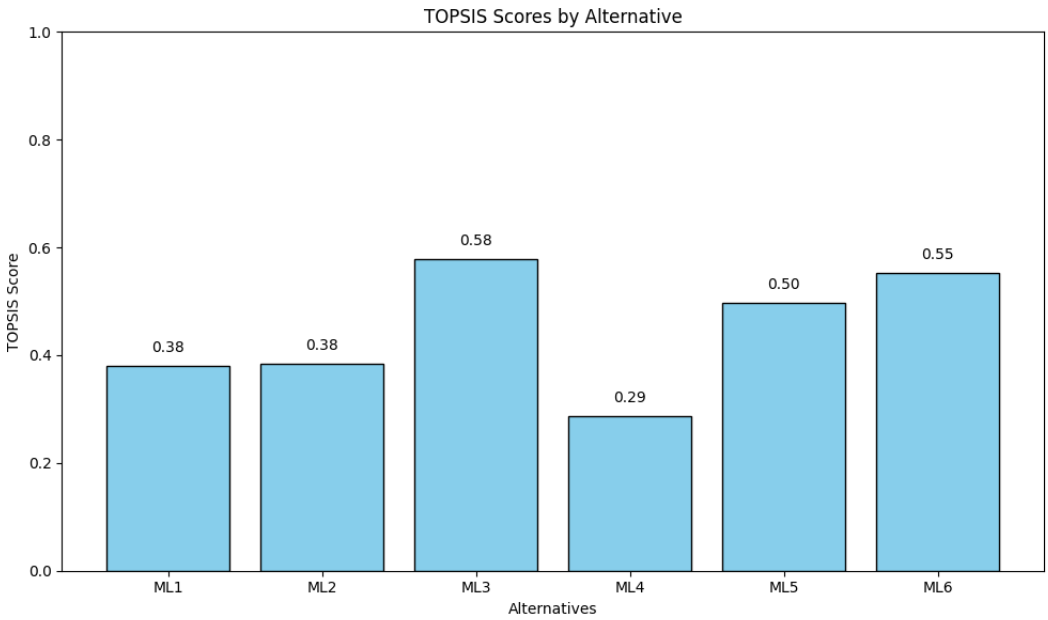

Figure 9: Bar chart (fava beans use case)

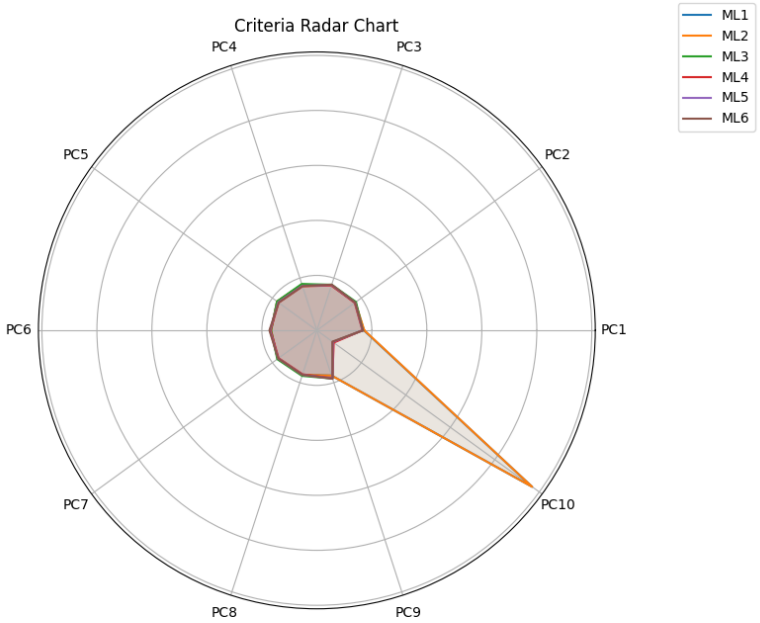

Figure 10: Radar chart (fava beans use case)

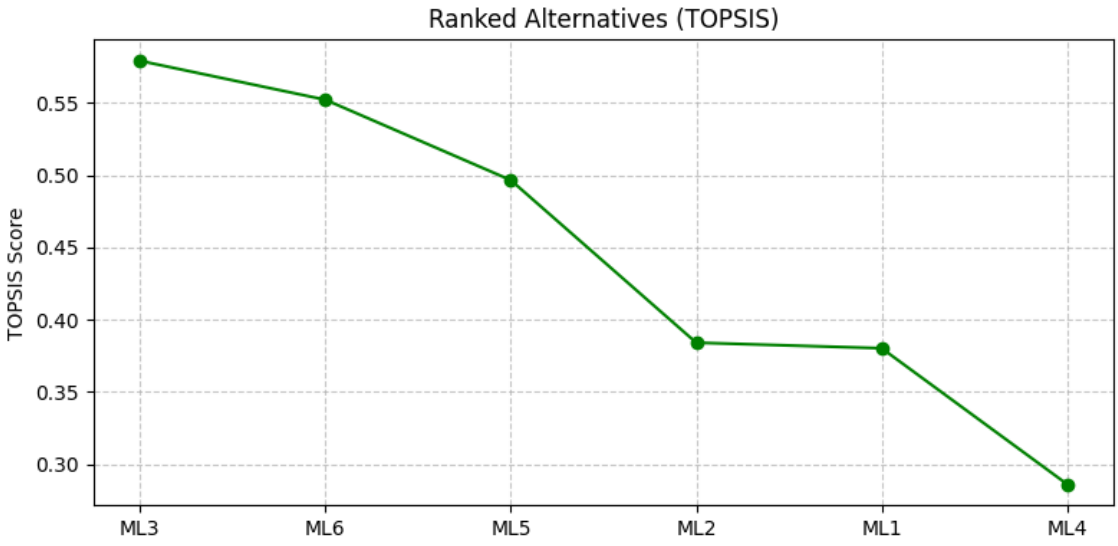

Figure 11: Ranked line chart (fava beans use case)

As can be observed ML3 is the best modeling pipeline in both scenarios (when criteria are not conflicted, when criteria are conflicted due to the addition of Loss performance criterion).

Evaluation in the Organic Honey Supply Chain

The second version of the TOPSIS algorithm is used for the evaluation of unsupervised anomaly detection models explored in the context of the ALLIANCE project for detecting fraud incidents in the organic honey supply chain. The data used for anomaly detection consists of records containing approximately 500 variables, and these variables might indicate potential anomalies based on specific characteristics of the honey, such as its origin or quality. The detection process involves the use of multiple anomaly detection models (e.g., CD[1], COF[2], COPOD[3], DEEP_SVDD[4], etc. which are part of the PyOD[5] library) used by The World Bee Project CIC (partner in the ALLIANCE project). These models work together to identify anomalous records that deviate from the normal patterns within the dataset. The data was processed using various characteristics such as species, family, order, date, sensor ID and readings. In addition, the dataset consists of multiple test runs of DNA and sensor data, with each dataset containing a combination of categorical and numerical data columns. The tests focus on detecting anomalies that could indicate fraud or other unusual activity in the honey supply chain. The real anomalies, if any, are used to test the models, while the fake anomalies are included in the training data to evaluate the effectiveness of the system. Anomaly scores are calculated based on the number of models that classify a dataset as anomalous. For example, a dataset with a score of 5 means that five of the 40 models have classified the dataset as anomalous.

In the case of this supply chain, there might be conflicting criteria when selecting the best performing models and parameters. For example, some models may flag records based on different variables and criteria that are characteristic of types of fraud or anomalies. Conflicts could arise when a certain type of anomaly is more indicative of a problem with a specific variable (e.g., an increase in one set of variables for honey from a specific region or decrease in another), making it necessary to evaluate these different characteristics in a multi-criteria decision-making approach such as TOPSIS. This situation requires balancing between multiple criteria such as true positive rate (i.e., correctly identified anomalies) and false positive rate (i.e., incorrectly flagged anomalies). As a result, the main goal is to maximize the true positive rate while minimizing the false positive rate.

To support multi-criteria evaluation of anomaly detection configurations in the organic honey supply chain, the following structured JSON request can be submitted to the TOPSIS service (see Table 9). The request includes seven alternatives (S1–S7), each representing a specific test scenario where various unsupervised machine learning models were used to detect potential fraud in DNA and sensor data. All scenarios applied a subset of models from a larger pool, including CD, COF, COPOD, DEEP_SVDD, DEVNET[6], DIFF[7], KPCA[8], and LOCI[9]. Scenarios S1, S2, and S3 applied the full model suite on DNA data with different clustering methods, while S4 and S5 used only three models (CD, COF, COPOD), with S5 adding KMeans over crop cover. Scenarios S6 and S7 tested the same full model suite on sensor data, with and without grouping by crop cover.

Each alternative was evaluated against four criteria:

[1] Clustering-based Local Outlier
[2] Connectivity-based Outlier Factor
[3] Copula-based Outlier Detection
[4] Deep Support Vector Data Description
[5] Python Outlier Detection
[6] Deep Anomaly Detection with Deviation Networks
[7] Deep Isolation Forest Feature Fusion
[8] Kernel Principal Component Analysis
[9] Local Correlation Integral

- PC1 (True Positive Rate): The ratio of correctly detected anomalies over total actual anomalies, a benefit criterion ("+") to be maximized.
- PC2 (False Positive Rate): The rate at which normal data points were wrongly flagged, a cost criterion ("-") to be minimized.
- PC3 (True Negative Rate): The accuracy of correctly identifying normal instances, a benefit criterion ("+") to be maximized.
- PC4 (False Negative Rate): The proportion of missed anomalies, also a cost criterion ("-") to be minimized.


```
{
 "matrix": [
  [0.8000, 0.0670, 0.9300, 0.0080],
  [1.0000, 0.0670, 0.9300, 0.0000],
  [1.0000, 0.3330, 0.6520, 0.0000],
  [0.8000, 0.3170, 0.6550, 0.0170],
  [0.2000, 0.0670, 0.9270, 0.0670],
  [0.6000, 0.0160, 0.9840, 0.0020],
  [0.5600, 0.0050, 0.9950, 0.0020]
 ],
 "weights": [0.30, 0.20, 0.30, 0.20],
 "signs": ["+","-","+","-"],
 "names_of_alternatives": ["S1", "S2", "S3", "S4", "S5", "S6", "S7"],
 "names_of_criteria": ["PC1", "PC2", "PC3", "PC4"]
}
```


Table 9: JSON request example (organic honey use case)

Table 10 presents the JSON response, whereas Figure 12, Figure 13 and Figure 14 presents the corresponding charts created by the TOPSIS service.


```
{
  "closeness_coefficient": [
    {
      "alternative": "S1",
      "score": 0.7968669076007607
    },
    {
      "alternative": "S2",
      "score": 0.8717665832163566
```

```
    },
    {
      "alternative": "S3",
      "score": 0.5
    },
    {
      "alternative": "S4",
      "score": 0.42573774688783156
    },
    {
      "alternative": "S5",
      "score": 0.4412336478857295
    },
    {
      "alternative": "S6",
      "score": 0.7390319850037169
    },
    {
      "alternative": "S7",
      "score": 0.7230835019650098
    }
  ],
  "execution_time_ms": 735.838,
  "message": "TOPSIS calculation completed successfully.",
  "radar_chart_path": "/charts/radar_chart_5928064f118e4849826601bd45430461.png",
  "bar_chart_path": "/charts/bar_chart_4cb223dd5d8d4dfdb5dcb5848af0120a.png",
  "ranked_line_chart_path": "/charts/ranked_line_chart_eabd27299de54e31be4eb8cfa7ccffb0.png"
}
```


Table 10: JSON response example (organic honey use case)

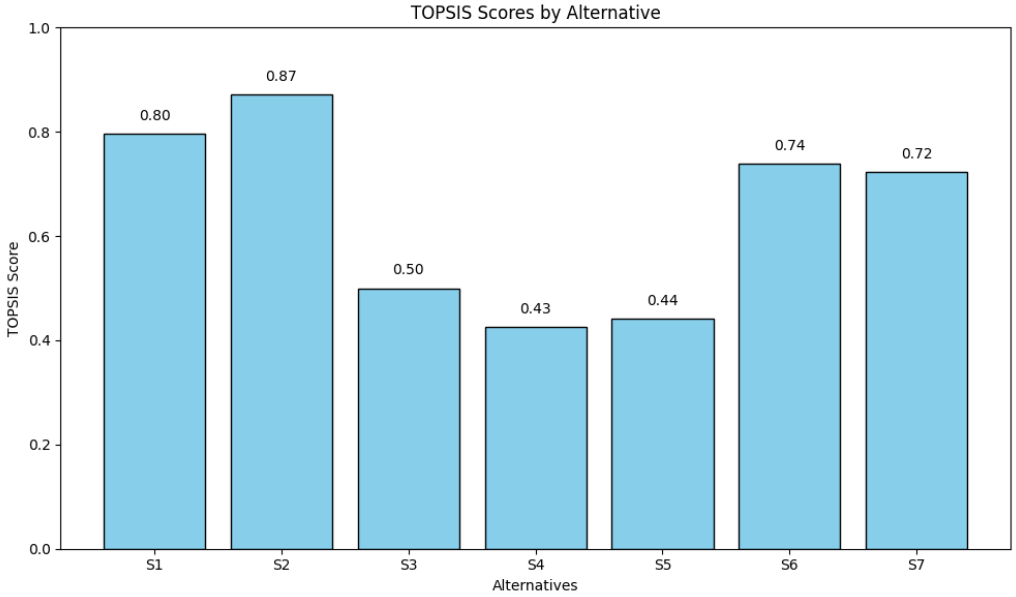

Figure 12: Bar chart (organic honey use case)

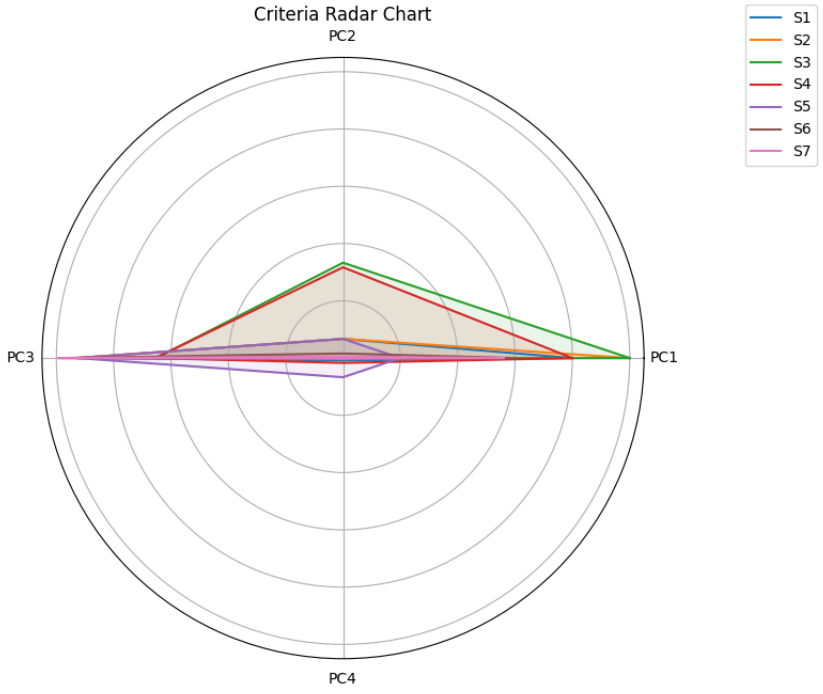

Figure 13: Radar chart (organic honey use case)

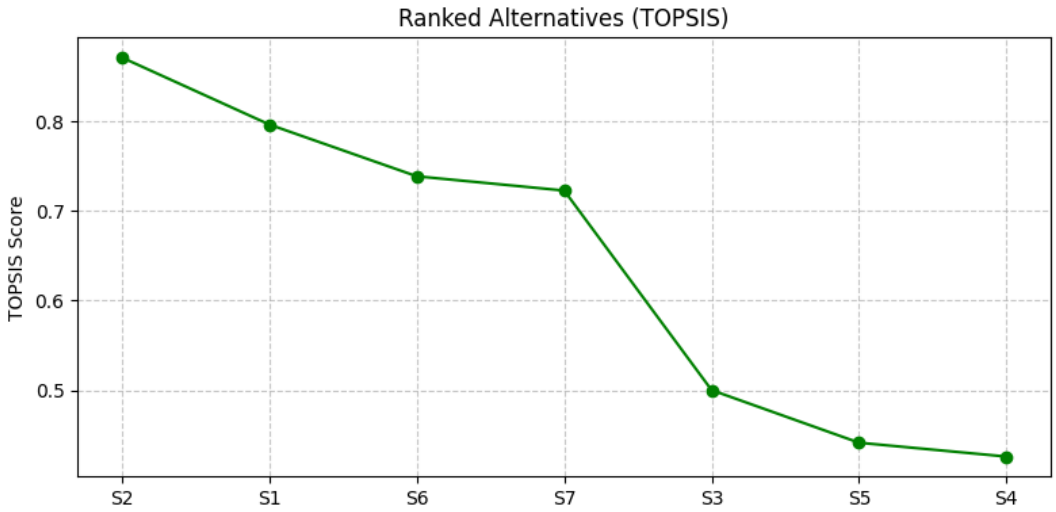

Figure 14: Ranked line chart (organic honey use case)

As it can be observed, scenario 2 seems to be the best one having an optimal trade-off between true positive rate and false positive rate.

5.3.2. Multi-criteria Decision Support for Supply Chain Stakeholders

As described in Sub-section 4.3 (*Multi-Criteria Decision Analysis Approach*) of Deliverable D3.2, the Analytic Hierarchy Process (AHP) [16] was recommended to support human-in-the-loop decision-making. In line with this, Netcompany-Intrasoft developed a decision-aid tool that enables users to express preferences between various alternatives and evaluation criteria. This tool is generic and flexible, designed for use by any stakeholder (in ALLIANCE) who can define a decision-making problem using a hierarchical structure comprising a goal, criteria, and alternatives. The primary objective is to rank the alternatives, assess the consistency of decision-makers, and guide users in answering three key questions:

- With respect to the overall goal, which criterion is the most important in the decision-making process?
- For a given criterion, which alternative is the most preferred?
- Based on the priorities of the criteria and the alternatives under each criterion, how are the alternatives scored and ranked overall?

To better illustrate the functionalities of the implemented decision support tool, we present two illustrative examples from the fava bean supply chain. These examples reflect real-world challenges where stakeholders need to weigh multiple factors to make informed, transparent, and consistent decisions. The first approach focuses on prioritizing types of fraud based on their potential impact and detection difficulties (see Subsection 4.3.2.1), while the second approach evaluates different mitigation strategies according to practical and economic considerations (see Subsection 4.3.2.2). Together, they show how the tool can support complex decision-making scenarios by structuring problems hierarchically and allowing users to systematically compare alternatives.

Prioritizing Fraud Risks

The main objective of this decision-making process is to identify and prioritize the most critical types of fraud in the fava bean supply chain. The assessment is based on the following criteria:

- C1: Difficulty of detection
- C2: Economic impact
- C3: Lack of consumer confidence (reputation risk)

The alternatives under consideration, representing different types of fraud, are:

- A1: Adulteration with cheaper beans (mainly from Bolivia)
- A2: Mislabelling – faba beans produced in neighbouring regions (mainly Galicia) and labelled as Asturias Faba Beans
- A3: False organic labelling

Figure 15 illustrates how users define the decision-making problem by setting the goal, establishing the criteria, and selecting the alternatives, thereby constructing the hierarchical structure.

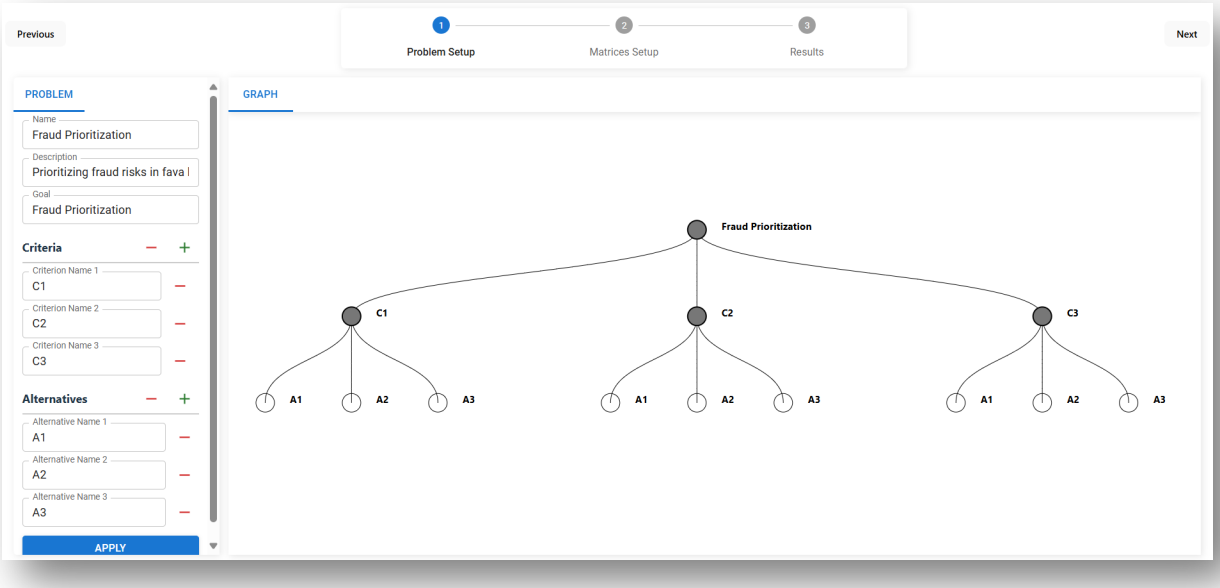

Figure 15: Problem definition (prioritizing fraud risks)

In the next step, once the goal is selected, a comparison matrix is created to allow the user to express preferences between criteria using the 1–9 scale described in Deliverable D3.2 (see Figure 16). Subsequently, by selecting a specific criterion, another comparison matrix is generated, enabling the user to express preferences between the alternatives (see Figure 17, Figure 18 and Figure 19).

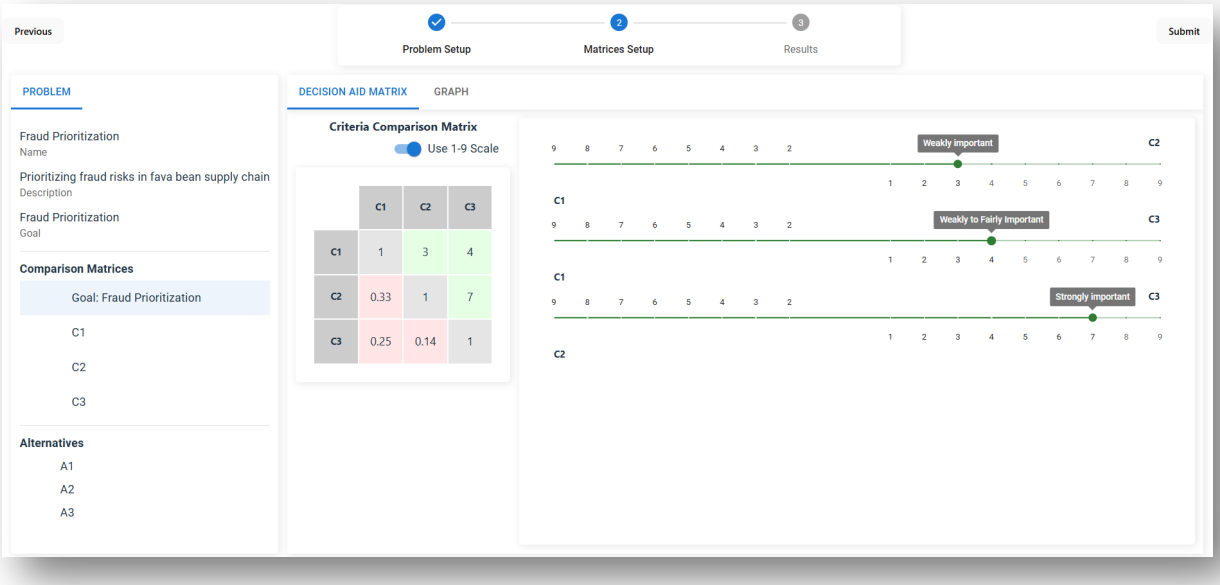

Figure 16: Criteria comparison (prioritizing fraud risks): C2 is weakly more important than C1; C3 is between weakly and fairly more important than C1; C3 is strongly more important than C2.

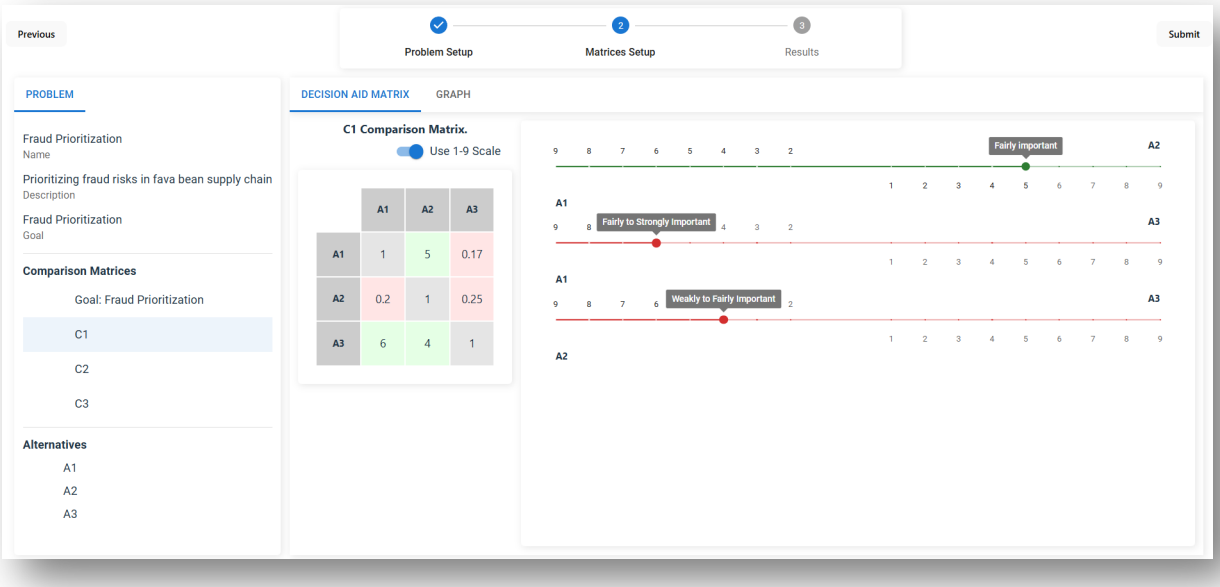

Figure 17: Comparison of alternatives with respect to C1 (prioritizing fraud risks): A2 is fairly more important than A1; A1 is between fairly and strongly more important than A3; A2 is between weakly and fairly more important than A3.

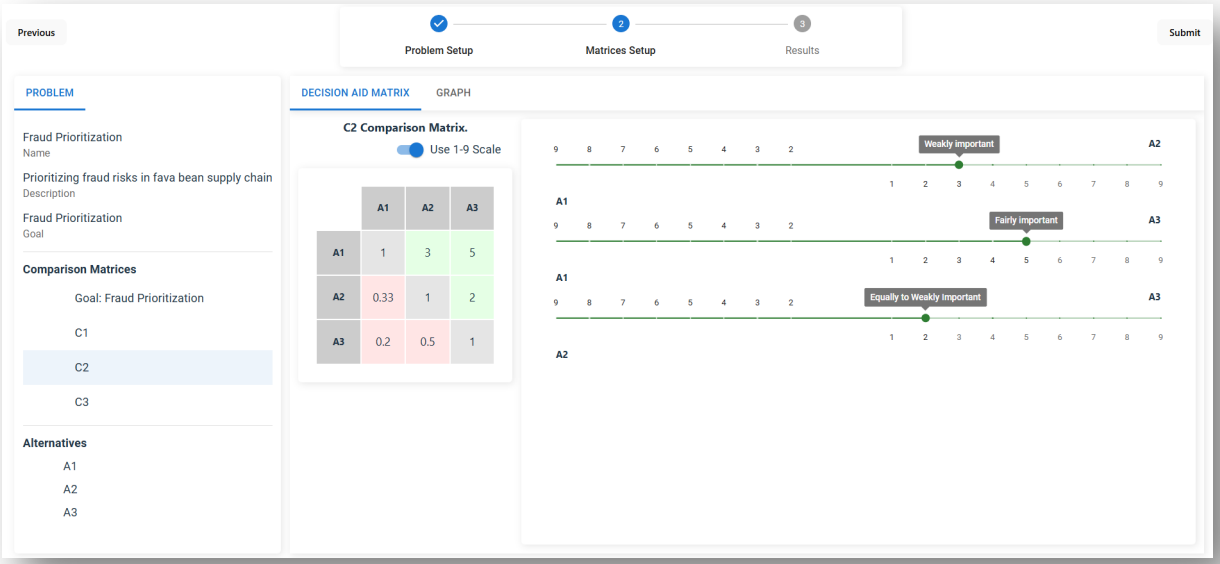

Figure 18: Comparison of alternatives with respect to C2 (prioritizing fraud risks): A2 is weakly more important than A1; A3 is fairly more important than A1; A3 Is between equally and weakly more important than A2.

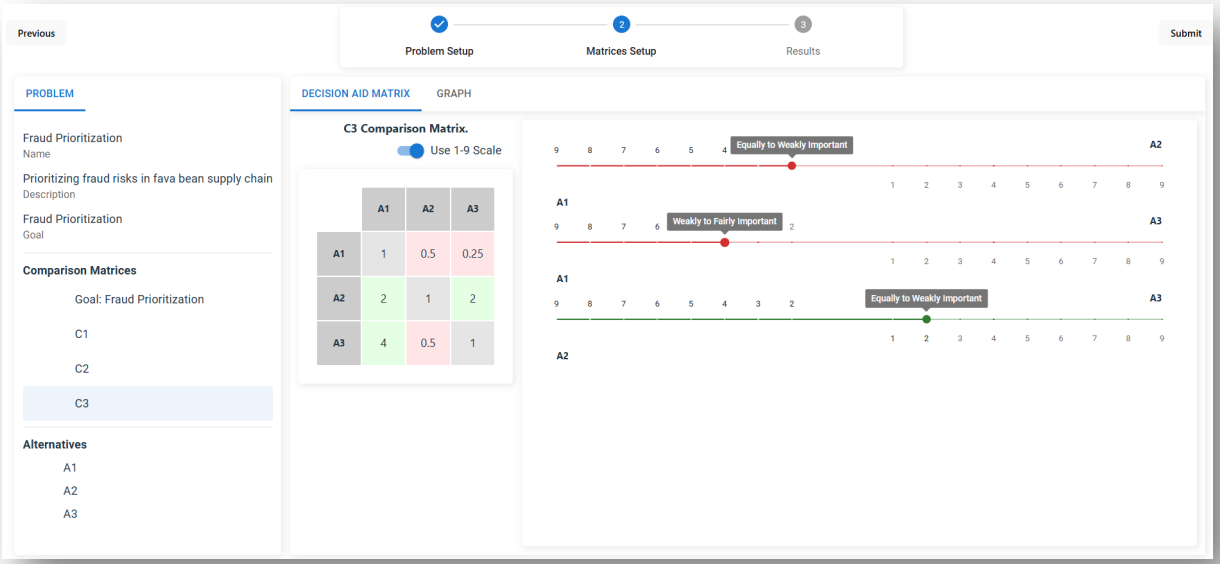

Figure 19: Comparison of alternatives with respect to C3 (prioritizing fraud risks): A1 is between equally and weakly more important than A2; A1 is between weakly and fairly more important than A3; A3 is between equally and weakly more important than A2.

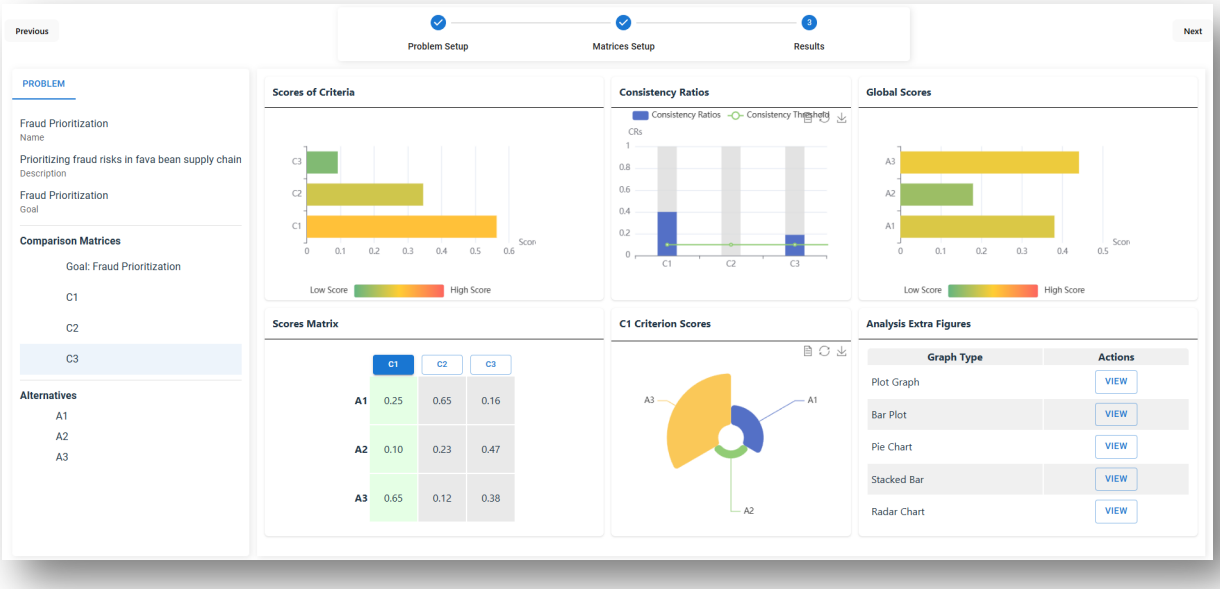

Figure 20: DSS results (prioritizing fraud risks)

Once the user has provided preferences for both the criteria and the alternatives under each criterion, the AHP proceeds to calculate the overall priorities. This involves synthesizing the input to determine the relative weights of the criteria and combining them with the alternative rankings to compute final scores for each alternative.

Given a pairwise comparison matrix $A = [a_{ij}]$, where a_{ij} represents the relative importance of element i over element j using 1-9 scale, normalization is carried out column-wise (i.e., sum each column of the matrix, and normalize each element by dividing it by the sum of its column). We then calculate the priority vector (weights) by averaging each row of the normalized matrix. Hence, the resulting vector $w = [w_1, w_2, \dots, w_n]$ represents the relative weights or priorities of the elements compared (i.e., alternatives or criteria). The final scores of alternatives in the AHP process are obtained by multiplying the weights of the criteria with the matrix of alternatives' weights per criterion.

To measure the consistency of the pairwise comparisons made by the decision maker, we can compute the Consistency Index (CI) as follows:

$$CI = \frac{\lambda_{max} - n}{n - 1} \qquad (9)$$

where λ_{max} is the largest eigenvalue of the pairwise comparison matrix and n is the number of criteria (or alternatives) being compared. The Consistency Ratio (CR) is calculated by dividing the CI by the RI (Random Consistency Index) which represents the consistency of randomly generated pairwise comparisons for matrices of different sizes according to the literature. If $CR \leq 10\%$, it means that the comparisons are consistent enough for the analysis to be valid.

In the final step (see Figure 20), the tool presents the results through a set of intuitive visualizations: a horizontal bar chart illustrating the scores of the criteria (highlighting the most influential one in the decision-making process); another horizontal bar chart showing the final scores and ranking of alternatives; and a vertical bar chart that displays CRs for each comparison matrix, helping to ensure decision coherence (ideally, consistency ratios should be below 10%). Additionally, a table summarizes the scores of alternatives per criterion,

complemented by a corresponding chart for quick visual reference. The tool also supports exporting supplementary visualizations to further assist users in interpreting the results.

Evaluating Mitigation Strategies

The objective of this decision-making process is to select the most appropriate fraud risk mitigation strategy for the fava bean supply chain. To achieve this, several criteria are considered, including the feasibility with existing resources (C1), the effectiveness of each strategy in combating fraud (C2), the cost-effectiveness of the solution (C3; both in terms of implementation and ongoing operations) and the time required for rapid implementation (C4). The alternatives evaluated include a data-driven decision support system for end-users responsible for quality controls (A1), blockchain-based traceability systems (A2), and training and awareness programs for growers and packers (A3).

It is important to note that the tool described is available in the context of ALLIANCE as a service: S.A.D.E - Swagger UI. It can be accessed using structured JSON data (see Table 11).


```
{
 "id": "680c59840afdc495ac7fa9bc",
 "created": "2025-04-26T03:56:52.111710Z",
 "updated": "2025-04-26T03:56:52.111742Z",
 "name": "Strategy Evaluation",
 "description": "Evaluating Mitigation Strategies",
 "goal": "Strategy Evaluation",
 "method": "normalized_column_sum",
 "uid": "",
 "alternatives": [
  {
   "name": "A1",
   "uid": "50a7ecb1-5155-4b9b-9b7a-c407905d01ab"
  },
  {
   "name": "A2",
   "uid": "0d5bb6a2-b04f-4f9c-aac7-cb88a4a0df18"
  },
  {
   "name": "A3",
   "uid": "563b3bb0-7c2c-48ca-bde1-a7da9cf9e2e6"
  }
 ],
 "criteria": [
```

```
  {
    "uid": "7124c9ea-c1f1-43d9-ad69-f009d8637502",
    "name": "C1",
    "parent": null
  },
  {
    "uid": "2b7df2e1-9d9b-4b28-b8c6-016d531ed07a",
    "name": "C2",
    "parent": null
  },
  {
    "uid": "6d441e01-a25c-4572-8d7e-245f97beb341",
    "name": "C3",
    "parent": null
  },
  {
    "uid": "b5e4da2b-a643-44d4-9c3c-e4743a844991",
    "name": "C4",
    "parent": null
  }
 ],
 "preferences": [
  {
    "hid": "criteria",
    "preferences": [
      [
        null,
        5,
        3,
        -3
      ],
      [
        null,
        null,
```

```
      5,
      -5
    ],
    [
      null,
      null,
      null,
      -3
    ]
  ]
 },
 {
  "hid": "C1-alternatives",
  "criterionUid": "7124c9ea-c1f1-43d9-ad69-f009d8637502",
  "preferences": [
    [
      null,
      4,
      4
    ],
    [
      null,
      null,
      3
    ]
  ]
 },
 {
  "hid": "C2-alternatives",
  "criterionUid": "2b7df2e1-9d9b-4b28-b8c6-016d531ed07a",
  "preferences": [
    [
      null,
      -4,
```

```
      3
    ],
    [
      null,
      null,
      -2
    ]
  ]
 },
 {
  "hid": "C3-alternatives",
  "criterionUid": "6d441e01-a25c-4572-8d7e-245f97beb341",
  "preferences": [
    [
      null,
      -4,
      3
    ],
    [
      null,
      null,
      3
    ]
  ]
 },
 {
  "hid": "C4-alternatives",
  "criterionUid": "b5e4da2b-a643-44d4-9c3c-e4743a844991",
  "preferences": [
    [
      null,
      -5,
      -6
    ],
```

```
    [
     null,
     null,
     2
    ]
   ]
  }
 ]
}
```


Table 11: Payload for the decision-making problem that aims to evaluate mitigation strategies

The "criteria" field in the payload defines the criteria used in the decision-making process. Each criterion has a unique uid (identifier) and a name. In addition to this, the "alternatives" field lists the alternatives (options) to be evaluated. Each alternative has a name and a uid.

The "preferences" section in the payload indicates the relative importance of the criteria compared to one another. This information is given in the form of matrices. For example, the matrix for hid: "criteria" shows how the criteria are compared. Furthermore, for each criterion, the preferences matrix shows how alternatives are compared with respect to that criterion (e.g., see hid: "C1-alternatives", "hid: C2-alternatives", etc.).

The following table (see Table 12) presents the structured JSON response of the service to the JSON request. In particular, the JSON response provides key insights into the evaluation results of the decision-making process. It includes the global scores for each alternative (A1, A2, A3), representing their overall rankings based on the criteria. Each criterion (C1, C2, C3, C4) has an individual score reflecting its importance in the decision-making process. In addition, the ratings of the alternatives per criterion show how each alternative performs under each specific criterion. The CR for the comparison matrix of the criteria, as well as the comparison matrices for the alternatives per criterion are provided. Finally, the answer contains links to graphics (e.g. bar plots, pie charts, radar charts) that visualize the evaluation results.


```
{
 "cr_criteria_matrix": 0.07393680280792869,
 "global_scores": {
  "A1": 0.28890600733072724,
  "A2": 0.4422192898435187,
  "A3": 0.268874702825754
 },
 "scores_of_criteria": {
  "C1": 0.27393151712071984,
  "C2": 0.16420894776305922,
  "C3": 0.08847788052986752,
```

```
  "C4": 0.4733816545863534
 },
 "scores_of_alternatives_per_criterion": [
  {
   "key": "C1_alternative_weights",
   "scores": [
    0.6387870709763411,
    0.24325224925024988,
    0.11796067977340885
   ]
  },
  {
   "key": "C2_alternative_weights",
   "scores": [
    0.3106666666666667,
    0.38966666666666666,
    0.2996666666666667
   ]
  },
  {
   "key": "C3_alternative_weights",
   "scores": [
    0.2581612258494337,
    0.6032644903397735,
    0.13857428381079281
   ]
  },
  {
   "key": "C4_alternative_weights",
   "scores": [
    0.08463845384871709,
    0.5454848383872042,
    0.3698767077640786
   ]
```

```
  }
 ],
 "crss": [
  {
   "key": "cr_C1",
   "cr": 0.11690590056727948
  },
  {
   "key": "cr_C2",
   "cr": 1.0613618774052846
  },
  {
   "key": "cr_C3",
   "cr": 0.18738066975485945
  },
  {
   "key": "cr_C4",
   "cr": 0.07393680280792869
  }
 ],
 "graphs": [
  {
   "name": "Plot Graph",
   "url": "/static/_basic_figure.png"
  },
  {
   "name": "Bar Plot",
   "url": "/static/_bar_plot_figure.png"
  },
  {
   "name": "Pie Chart",
   "url": "/static/_pie_chart_figure.png"
  },
  {
```

```
    "name": "Stacked Bar",
    "url": "/static/_stacked_bar_chart_figure.png"
   },
   {
    "name": "Radar Chart",
    "url": "/static/_radar_chart_figure.png"
   }
 ]
}
```


Table 12. JSON response for the decision-making problem that aims to evaluate mitigation strategies

6. Vulnerability Risk Assessment

6.1. Overview

The Vulnerability Risk Assessment is a combination of two different processes that complement each other. Vulnerability Assessment refers to a systematic examination of the food value chain to identify security deficiencies and weaknesses that could be exploited, while the main objective of the Risk Assessment is to identify and prioritize security risks by analyzing threats and vulnerabilities—thereby informing which controls or responses should be implemented. In ALLIANCE, we develop a **Vulnerability Risk Assessment Management Framework (VRAMF)**, which is our basis for evaluate the threats and the food fraud possibilities in the developed FSCs. VRAMF is used for the detection of the most critical control points in the FSCs, which are the points that extra control is needed for detecting and deteriorating food fraud.

To detect potential food fraud incidences at the earliest possible stage, our approach implemented a targeted sampling strategy at the Critical Control Points (CCPs) within the FSCs. This effort relied on and built directly from on the outcomes produced in T2.1, where we developed a comprehensive understanding of the FSCs operations and flow of goods through surveying key stakeholders. We relied on the Delphi technique to refine our data collection. Multiple questionnaire rounds were organized to collect the necessary information from the stakeholders. Each survey was dynamically tailored to the previous round responses. When questionnaire results required further clarifications, we organized additional follow-up meetings/calls with the relevant stakeholders. This iterative process allowed for gathering precise stakeholder-driven information about each step in the FSCs flow. The purpose of this effort, apart from understanding the information flow of the FSCs that was the main goal of T2.1, was to identify the CCPs of the FSCs and introduce quality controls at these points, in order to strengthen our ability to monitor and validate critical procedures.

After defining the initial set of CCPs for each FSC (with each CCP indicating the location/point within the FSC where samples are collected and quality control results are generated, we also had the first set of data that fed the EWDSS App. The purpose of facilitating the FSC admin taking decisions against potential food frauds. The assessment of the performance of the EWDSS App is an ongoing continuous process. We iteratively assess the capabilities of EWDSS to detect food frauds and identify opportunities to propose additional CCPs (e.g. new sampling points and quality control data) when deemed necessary. When an introduction of new CCP is evaluated and considered necessary to improve significantly the EWDSS detection accuracy, the CCP is appended to the existing set and the VRAMF process resumes interacting with EWDSS. Figure 21 illustrates this dynamic relationship between the VRAMF and the EWDSS highlighting the impact on the definition of the CCPs of each FSC and showing how each iteration refines and expands the CCPs definition.

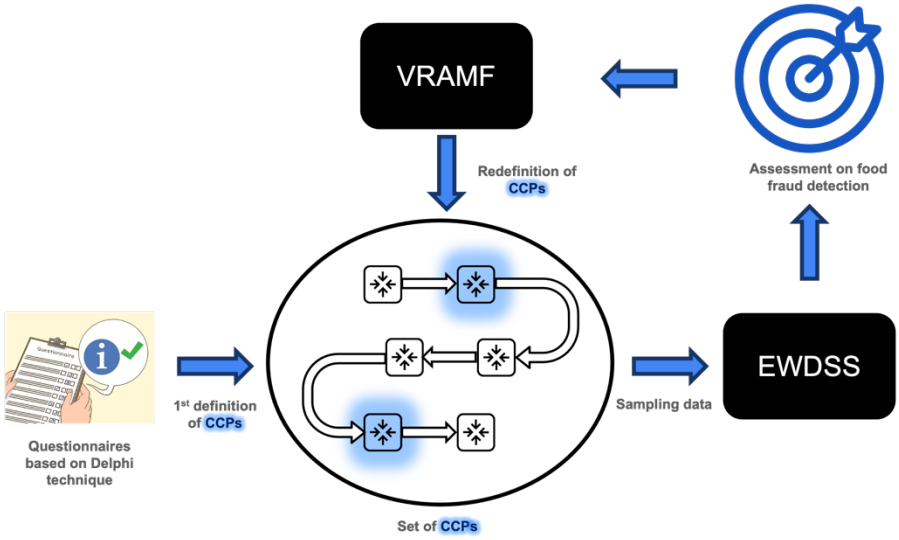

Figure 21: VRAMF - EWDSS interaction and CCPs definition.

The integration of the Critical Control Points (CCP) with the Blockchain and the Early Warning Decision Support System (EWDSS) is foreseen as an advanced-level feature, aimed at enhancing cross-platform coordination, security, and data traceability. The necessary architectural design and interfacing mechanisms have been carefully designed, ensuring that the integration can be effectively realized with minimal effort. The architectural design provides the integration points, data flow models, and interaction protocols between CCP, Blockchain, and EWDSS which has been fully defined and validated.

6.2. CCPs on the Feta Cheese chain

For the Feta Cheese food supply chain, guided by the results produced in T2.1, we introduced 2 CCPs:

- **Milk delivery**: The point where the milk is delivered by the Truck Driver to the Reception Manager, and
- **Post-Pasteurization**: The point where the Production Manager receives the milk, after pasteurization, for cheese production.

As it is depicted in the left side of Figure 22, these CCPs show the locations within FSC where samples are collected and forwarded to the Quality Control manager, who is responsible for conducting the quality control tests. The resulting data is fed into the EWDSS. The EWDSS, in turn, analyses those data to recommend appropriate actions and support the decision-making process for the FSC administrator in mitigating possible food frauds. As mentioned above, the VRAMF continuously evaluates the performance efficiency of the EWDSS and determines whether an additional set of CCPs should be introduced.

For example, for the Feta Cheese food supply chain, that was the first developed among all FSCs, our analysis indicated that adding an additional sampling point would further enhance fraud detection, which is:

- **Pasteurization:** the point where the Pasteurization Manager receives the milk from the Reception Manager.

The whole set of CCPs are shown on the right side of Figure 22.

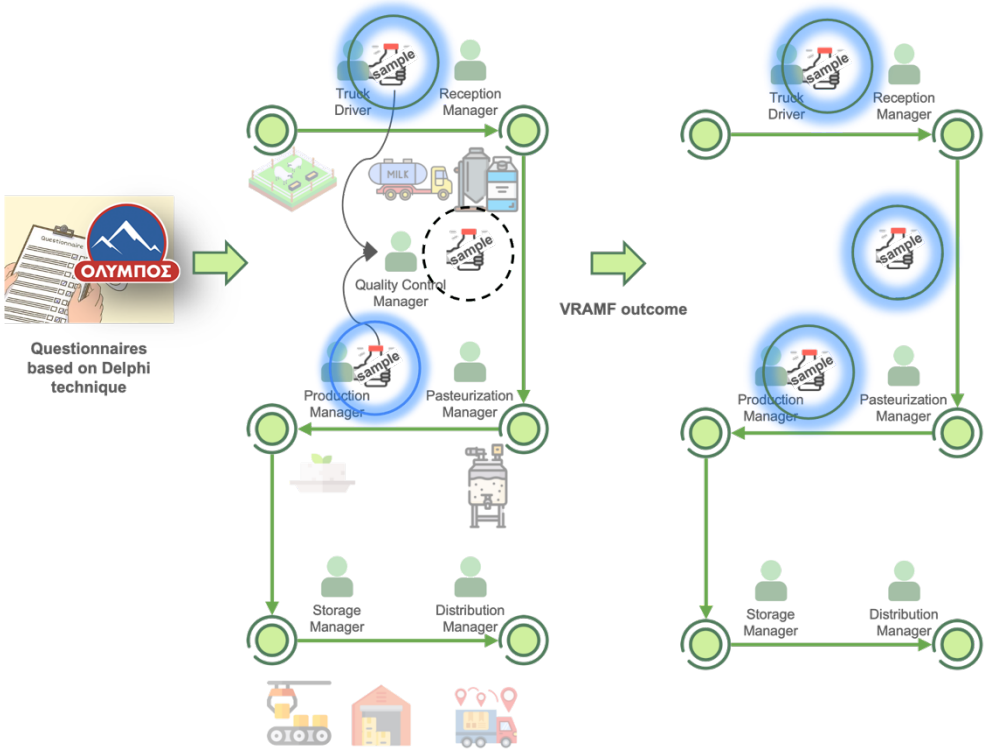

Figure 22: CCPs on the Feta Cheese chain.

6.3. CCPs on the Olive Oil chain

For the Olive Oil supply chain, based on the results of T2.1 and the analysis of the VRAMF process, we identified 2 CCPs:

- **Olive Reception:** the point where the harvested olive fruits are delivered to the Reception Manager, and
- **Milling handoff**: the point where the olive fruits are received by the Milling Manager for processing.

The complete set of CCPs for this chain, is depicted in Figure 23.

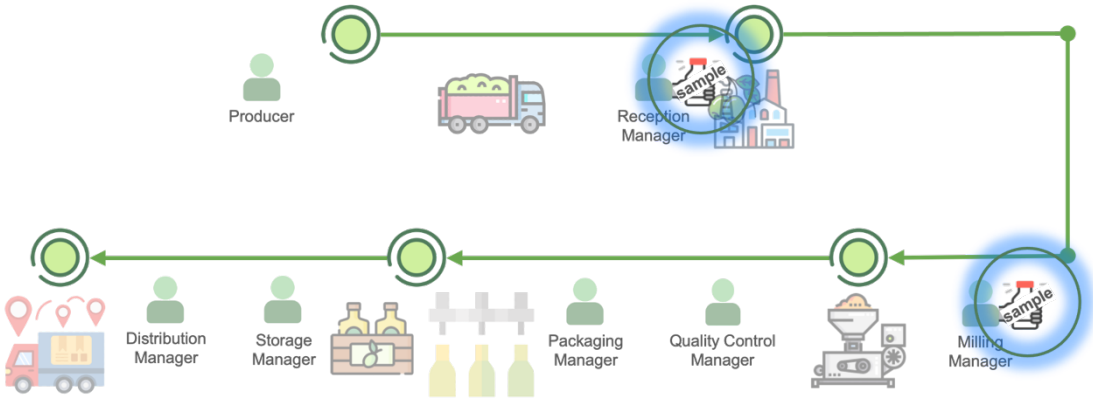

Figure 23: CCPs on the Olive Oil chain.

6.4. CCPs on the Organic Honey chain

For the Organic Honey supply chain, based on the results of T2.1 and the analysis of the VRAMF process, we identified 1 CCP:

- **Honey Reception:** The point where the honey is delivered to the Association Manager.

This CCP is shown in Figure 24.

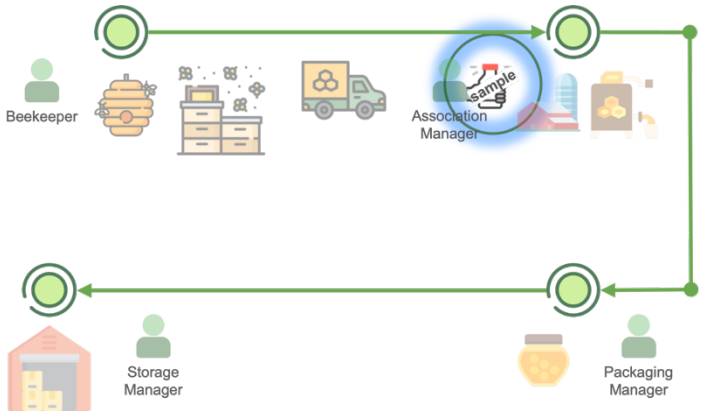

Figure 24: CCPs on the Organic Honey chain.

6.5. CCPs on the Faba Beans chain

For the Faba Beans supply chain, based on the results of T2.1 and the analysis of the VRAMF process, we identified 1 CCP:

- **Faba Reception**: the point where the Faba Beans are delivered to the Packaging Manager.

This CCP is visualized in Figure 25.

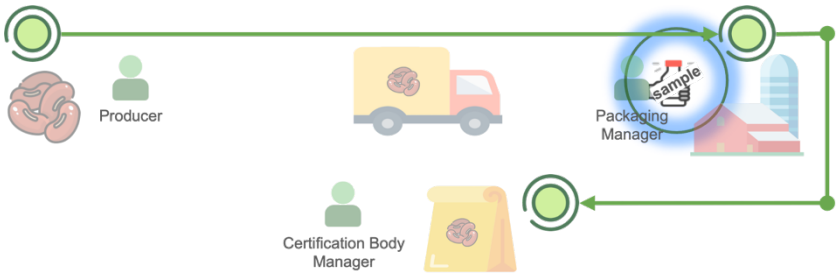

Figure 25: CCPs on the Faba Beans chain.

6.6. CCPs on the Lika Potatoes chain

For the Lika Potatoes supply chain, based on the results of T2.1 and the analysis of the VRAMF process, we identified 2 CCPs:

- **Potatoes Reception**: the point where the Lika Potatoes are delivered to the Reception Manager, and
- **Potatoes Pre-Packaging**: the point where the Lika Potatoes are delivered the Packaging Manager.

These CCPs are visualized in Figure 26.

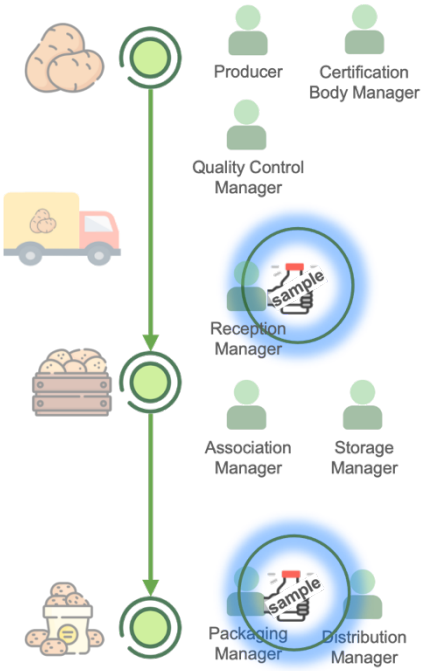

Figure 26: CCPs on the Lika Potatoes chain.

6.7. CCPs on the Organic Pasta chain

For the Organic Pasta chain, based on the results of T2.1 and the analysis of the VRAMF process, we identified 2 CCPs:

- **Wheat delivery**: the point where the wheat is delivered to the Wheat Reception Manager, and
- **Semolina delivery:** the point where the produced semolina is delivered to the Semolina Reception Manager.

These CCPs are visualized in Figure 27.

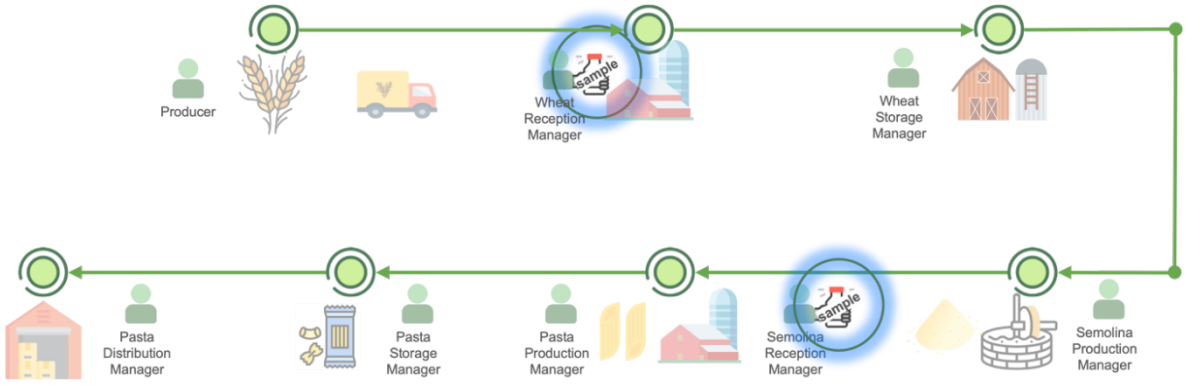

Figure 27: CCPs on the Organic Pasta chain.

6.8. CCPs on the Arijle Raspberry chain

For the Arijle Raspberry chain, based on the results of T2.1 and the analysis of the VRAMF process, we identified 2 CCPs:

- **Arijle Raspberries delivery:** the point where the Arijle Raspberries are delivered to the Reception Manager, and
- **Packaging**: the point where frozen raspberries are delivered to the Packaging Manager for packaging.

These CCPs are shown in Figure 28.

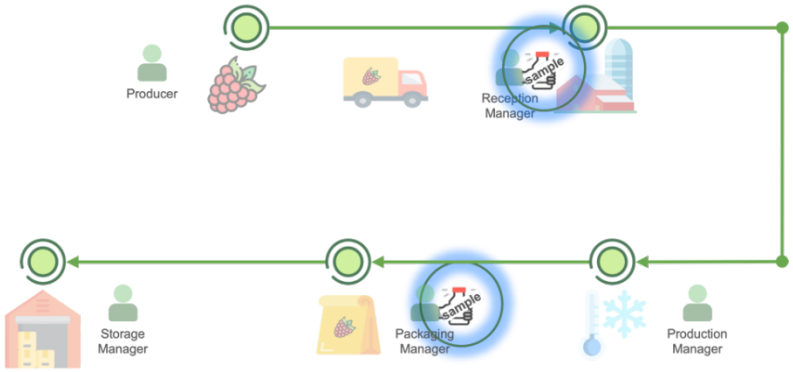

Figure 28: CCPs on the Arijle Raspberry chain.

7. Interoperability between Food Supply Chains

7.1. Overview

The interoperability of the data exchanged on the FSCs is a crucial process that focuses on the harmonization of the heterogeneities between these data. The idea is that the datastores (Blockchain and Off-chain) use the **GS1 EPCIS** (Electronic Product Code Information Services) standard and its companion **CBV** (Core Business Vocabulary) to define the fields of the tables in these datastores. The utilization of these standards for the naming of the table fields facilitates easy data sharing within organizations and stakeholders across the entire FSC.

CBV is designed to facilitate interoperability in EPCIS data exchange by providing standard values for vocabulary elements to be included in EPCIS data. The standard recognizes that the greatest interoperability is achieved when all data conforms to the standard and recognizes that their users may need to extend the standard in certain situations. To that end, this standard defines two levels of conformance for EPCIS documents:

- CBV-Compliant: An EPCIS document that ONLY uses vocabulary identifiers specified in the CBV standard in the standard fields of EPCIS events.
- CBV-Compatible: An EPCIS document that uses a COMBINATION of vocabulary identifiers specified in the CBV standard and other identifiers that are outside the standard.

Our focus is on facilitating data sharing, when the **products are moved from the producer to the retailer, using a CBV-Compatible EPCIS description of the events**. In this way, interoperability is facilitated, and data can be exchanged between the FSCs, since all kinds of containers (packages, pallets or boxes) are presented in a uniform manner, regardless of the food products they contain. Thus, the retailer is capable of equitably managing all packages or boxes.

In terms of T2.5, we:

- collected visibility goals and requirements from MASOUTIS and MIGROS,
- documented the FSC flows,
- broke each FSC flow into a series of discrete business steps,
- decided which business steps require visibility events,
- modeled the completion of each step as a visibility event - Understand what information is needed from a business application's perspective,
- decided what data fields are to be included in the visibility event and
- determined the vocabularies that populate each data field according to the CBV standard.

Currently, we use events specified by GS1 EPCIS, especially for the description of the data that are related to the packaging and transferring of the products to the retailer, which are generated at the V1, V2 and V3 points of the simple business process depicted in Figure 29.

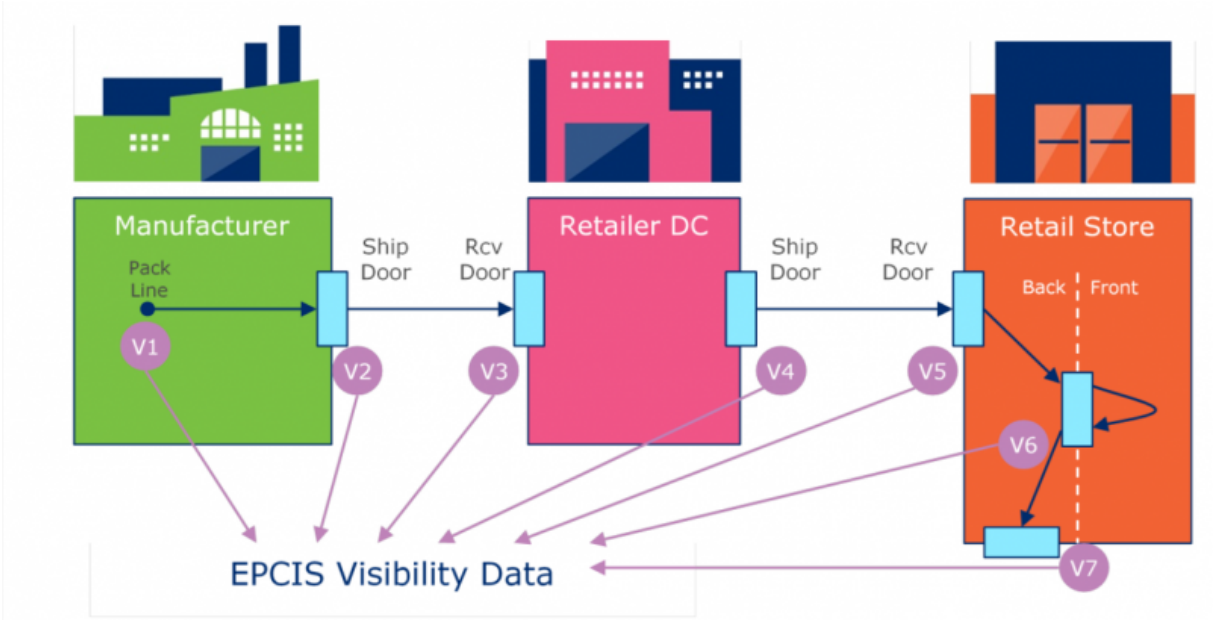

Figure 29: EPCIS visibility data during a simple business process.

Each of these events can be:

- **ObjectEvent**: which happens to one or more food products, when they are e.g. shipped or received. This is the simplest and most commonly used type of event. Instance-level EPC (Electronic Product Code) or class-level EPC without any relationship may appear in the ObjectEvent.
    - Instance-level EPCs: can be used to assign a unique number to differentiate each product
    - Class-level EPCs: If multiple objects are associated with the same identifier, then these types of identifiers may be considered class-level identifiers.

  A simple example is a pallet that is shipped or received using the pallet's code.

- **AggregationEvent**: which happens to multiple food products that are physically aggregated together or disaggregated from each other. The AggregationEvent is reversible - meaning that upon the disaggregation, original objects can be obtained. For example, aggregating cases onto a pallet, or removing cases from a pallet. This is the next most common type of event after ObjectEvent, and these two event types together cover most events in a typical business process.

7.2. GS1 EPCIS events from the Feta Cheese chain

In the following example, we showcase some GS1 EPCIS events used in the Feta Cheese chain. In the packaging stage, the Feta Cheese packages are first packaged into boxes and then several boxes are packaged into a pallet with a pallet code.

7.2.1. GS1 EPCIS Event for Feta Cheese packages

Each Feta Cheese product is a package, containing a specified quantity and quality of Feta Cheese. In general, Feta Cheese can be shipped in bulk, but in our case, each product is a

package. The creation of each package generates an **ObjectEvent**, which is shown in Figure 30, where:

- **Event type** is of type *Object*,
- **EPCs** list contains the ID of the created Feta Cheese package (ID: 0),
- **Event Time** is the time that this is event took place,
- **Record Time** is the time that this EPCIS description was generated,
- **Read Point** is an SGLN ID that corresponds to the location that the creation of the Feta Cheese package took place. SGLN is an extension of GLN that stands for 'Global Location Number'.
- **Business Step** is *commissioning*,
- **Disposition** is *active*.

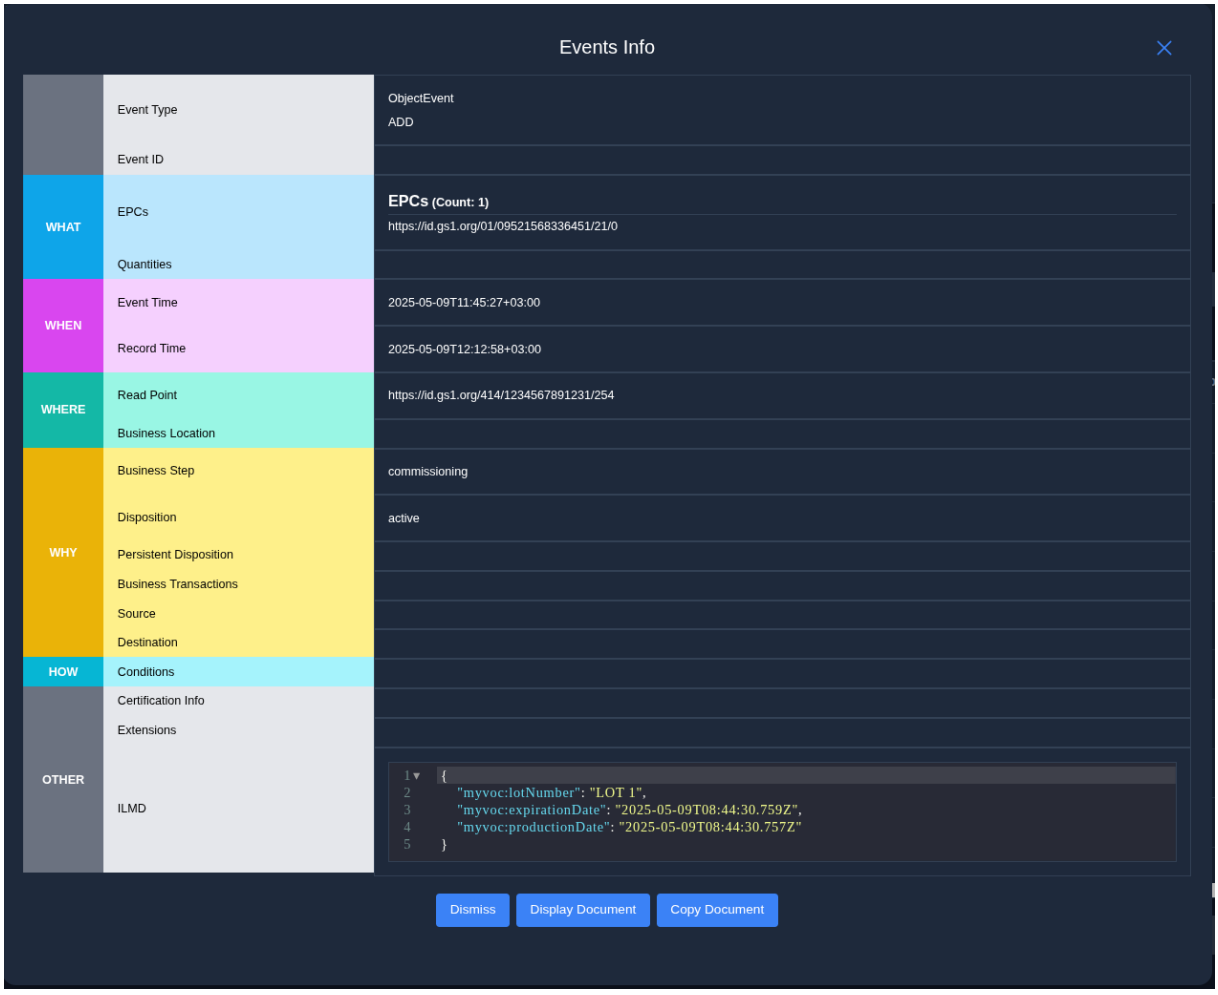

Figure 30: GS1 EPCIS Aggregation Event presenting the creation of a Feta Cheese package.

As we have already explained in previous deliverable D2.3, the EPCIS events are structured in a way that they clearly answer at least the four questions that are depicted in Figure 30, which are the following:

- **WHAT** generated this event?
    - The creation of a new object (*Event Type*) with *EPC* (of the Feta cheese package) equal to 0.
    - *Quantities* is not used, since the product has a serialized EPC, which means that it is shipped as a package and not in bulk.
- **WHEN** was this event generated?
    - *Event Time* gives the date and time at which the creation of the Feta Cheese package took place.
    - *Record Time* is the date and time that this event is stored in the EPCIS repository and does not provide information concerning the business step of the FSC.
- **WHERE** was this event created?
    - *Read Point* is the exact location, identified by SGLN, where the creation of the Feta Cheese package took place, within the OLYMPOS factory.
    - *Business Location* is the position where this package will be placed after this event, which is not defined at this moment, since its placement in warehouse units happens in following events.
- **WHY** is this event generated?
    - *Business Step* indicates what type of process was taking place at the time of the event within the context of the FSC process. Some examples from CBV include 'commissioning', 'receiving', 'picking', 'shipping' or 'packaging'. Obviously, commissioning is the business step during which the package is associated with a new EPC.
    - *Disposition* represents the state of the product immediately after the current EPCIS event. Some vocabularies from CBV for the disposition are 'expired', 'recalled', 'active', 'in_transit' or 'in_progress'. In this event, disposition is populated as 'active', which indicates that the package is ready for additional processing.

7.2.2. GS1 EPCIS Event for Feta Cheese boxing

The packaging of Feta Cheese packages into boxes is an **AggregationEvent** since multiple Feta Cheese packages are aggregated together to create a box. The generated GS1 EPCIS event is shown in Figure 31, where:

- **Event type** is of type *Aggregation*,
- **Parent ID** is the ID of the box (101),
- **CHILD EPCs** list contains the IDs of the Feta Cheese packages (0 to 9 are included in this box),
- **Event Time** is the time that this is event took place,
- **Record Time** is the time that this EPCIS description was generated,

- **Read Point** is the SGLN ID that corresponds to the location that boxing of Feta Cheese packages took place.
- **Business Step** is *packing*,
- **Disposition** is *in_progress*.

Two more boxes with IDs 102 and 103 have been created, as we will see in the following Aggregation Event that all boxes are put together in a pallet.

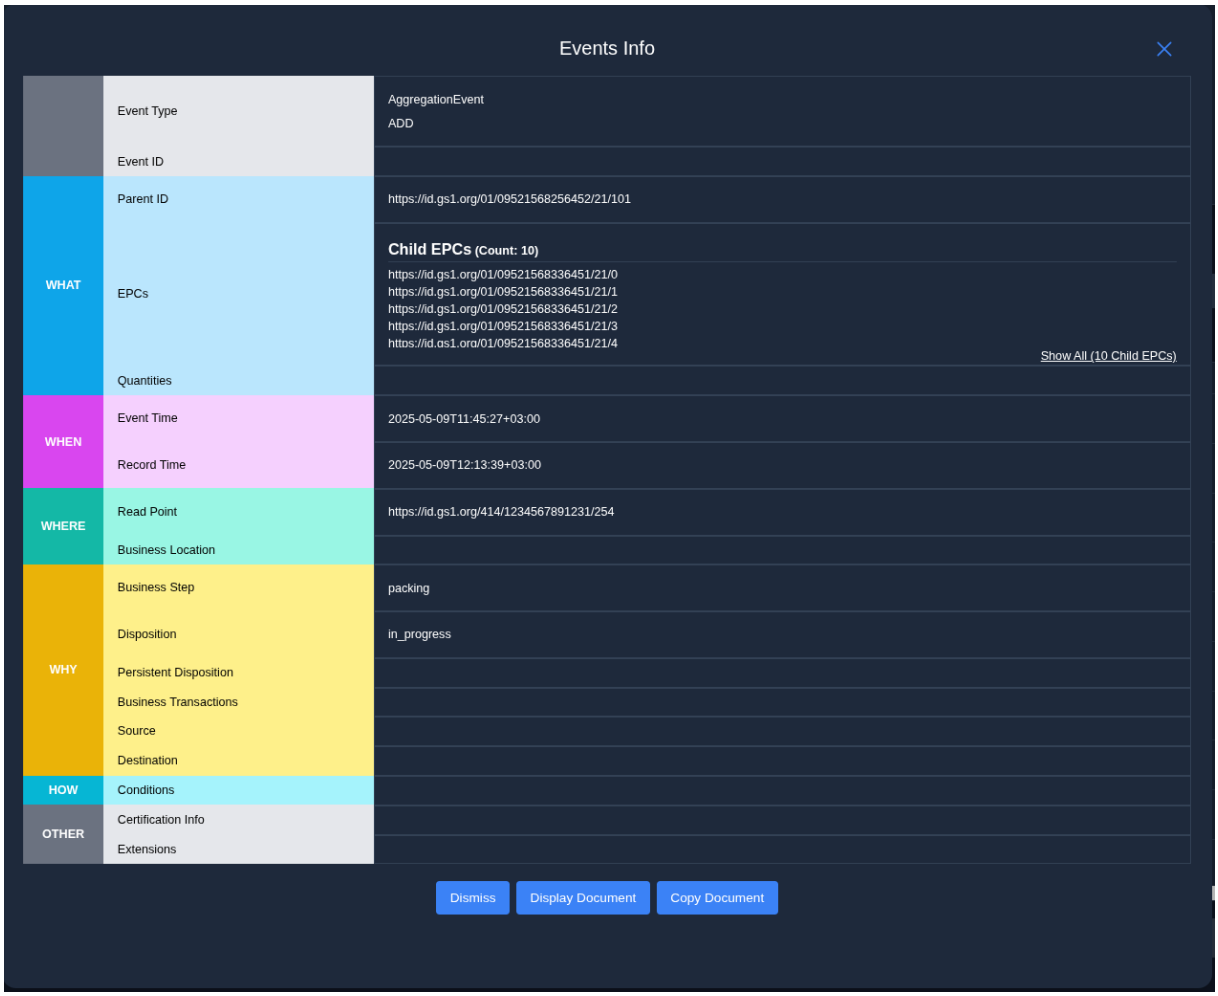

Figure 31: GS1 EPCIS Aggregation Event presenting a box of Feta Cheese packages.

This EPCIS event answers again the four questions that are depicted in Figure 32, which are the following:

- **WHAT** generated this event?
  - The aggregation (*Event Type*) of 5 child *EPCs* (Feta cheese packages) to 1 box identified by the *Parent ID*.
  - *Quantities* is not used, since there are serialized EPCs (packages)
- **WHEN** was this event generated?

- *Event Time* gives the date and time at which the boxing of the Feta Cheese packages took place.
- *Record Time* is the date and time that this event is stored in the EPCIS repository and does not provide information concerning the business step of the FSC.

- **WHERE** was this event created?
  - *Read Point* is the exact location, identified by SGLN, where the boxing of the Feta Cheese packages took place, within the OLYMPOS factory.
  - *Business Location* is the position where boxes will be placed after this event, which is not defined at this moment, since their placement in warehouse units happens in following events.
- **WHY** is this event generated?
  - *Business Step* indicates that the type of process at the time of the event is 'packaging'. Obviously, packaging is the business step during the boxing of the Feta Cheese packages.
  - *Disposition* is populated as 'in_progress', which indicates that the packages are moving normally through the supply chain and are not recalled.

7.2.3. GS1 EPCIS Event for Feta Cheese pallets

The next step is to package the boxes into a pallet to ship them to the retailer. Once again, an **AggregationEvent** is generated, since several boxes are aggregated together to create a pallet. The generated event is depicted in Figure 32. Here we can see that:

- **Event type** is of type *Aggregation*,
- **Parent ID** is the ID of the pallet,
- **CHILD EPCs** list contains the IDs of the Feta Cheese boxes (101 to 103 are included in this pallet),
- **Event Time** is the time that this is event took place,
- **Record Time** is the time that this EPCIS description was generated,
- **Read Point** is an SGLN ID that corresponds to the location that pallet was loaded,
- **Business Step** is *packing***,**
- **Disposition** is *in_progress***.**

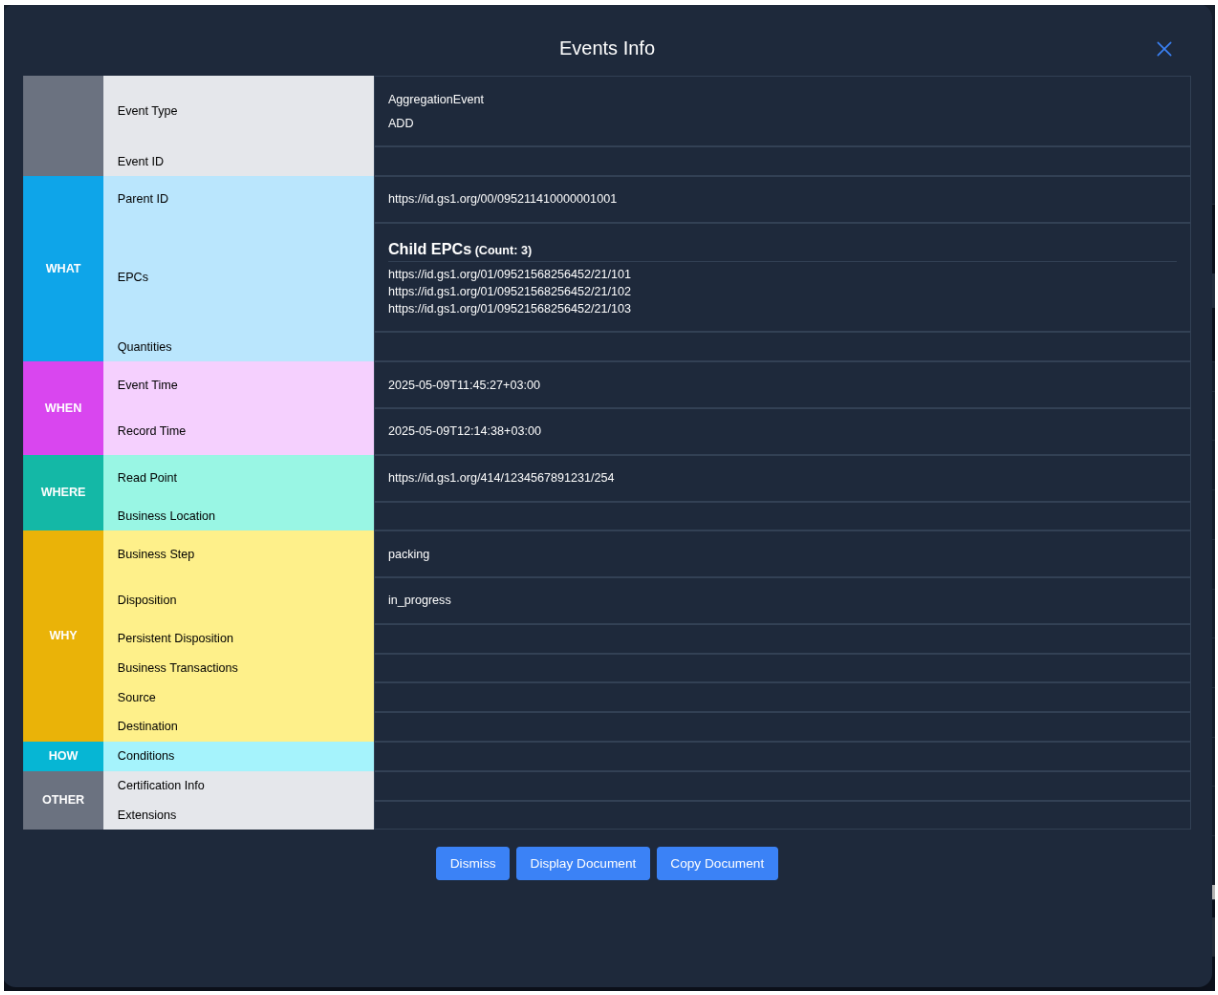

Figure 32: GS1 EPCIS Aggregation Event presenting a pallet of Feta Cheese boxes.

The values of these fields provide again the answers to the four questions: what, when, where and why.

8. Conclusion

This deliverable *D2.4 -Final AI-enabled tools for Vulnerability Risk Assessment, Early Warning Indication and Decision Support Preventive Actions* concludes the extensive work carried out under WP2 to M30, documenting the development, and implementation of key technological components to improve risk awareness and cross-system communication in quality-labelled food supply chains. It provides a detailed presentation of the developed services underpinning the Early Warning and Decision Support System, the Vulnerability Risk Assessment tool, and the Interoperability System built on GS1 and EPCIS standards. Each system was analysed and reported on, highlighting its specific functions, configuration, and role within the broader ALLIANCE architecture.

EWDSS is designed to facilitate the timely detection of anomalies and emerging threats, as well as to enable stakeholders to take preventive actions. The Vulnerability Risk Assessment component provides structured insights into potential weaknesses along the supply chain, enabling proactive risk management. While, the Interoperability System ensures seamless and standardised data exchange between heterogeneous systems and stakeholders, enabling integrated traceability and compliance verification.

All three systems have a high degree of operational maturity and are prepared for extensive testing and validation during the pilot phases. Their combined deployment marks a crucial step towards the realisation of a responsive digital infrastructure tailored to the agri-food value chains. The results of this deliverable form the basis for the final integration of these tools into the ALLIANCE platform and their deployment in the pilot scenarios.

REFERENCES

1. “GS1 standards repository.” Accessed: Apr. 29, 2025. [Online]. Available: https://ref.gs1.org/standards/#epcis
2. “GS1 standards repository.” Accessed: Apr. 29, 2025. [Online]. Available: https://ref.gs1.org/standards/#cbv
3. Z. Zheng, S. Xie, H. N. Dai, X. Chen, and H. Wang, “Blockchain challenges and opportunities: A survey,” *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352–375, 2018, doi: 10.1504/IJWGS.2018.095647.
4. T. Hepp, M. Sharinghousen, P. Ehret, A. Schoenhals, and B. Gipp, “On-chain vs. off-chain storage for supply-and Blockchain integration,” *IT - Information Technology*, vol. 60, no. 5, pp. 283–291, 2021, doi: 10.1515/ITIT-2018-0019.
5. Y. Xu *et al.*, “Artificial intelligence: A powerful paradigm for scientific research,” *The Innovation*, vol. 2, no. 4, p. 100179, Nov. 2021, doi: 10.1016/J.XINN.2021.100179.
6. K. L. Hulebak and W. Schlosser, “Hazard analysis and critical control point (HACCP) history and conceptual overview,” *Risk Anal*, vol. 22, no. 3, pp. 547–552, 2002, doi: 10.1111/0272-4332.00038.
7. “DELPHI PROCESS: A METHODOLOGY USED FOR THE ELICITATION OF OPINIONS OF EXPERTS.” Accessed: Apr. 29, 2024. [Online]. Available: https://apps.dtic.mil/sti/citations/AD0675981
8. K. Green, J. Armstrong, and A. Graefe, “Methods to elicit forecasts from groups: Delphi and prediction markets compared.,” *Foresight: The International Journal of Applied Forecasting*, vol. 8, pp. 17–20, Dec. 2008, doi: 10.2139/ssrn.1153124.
9. J. J. Thakkar, “Multi Criteria Decision Making, First Edition,” *Springer International Publishing: New York*, 2021, Accessed: Apr. 29, 2024. [Online]. Available: https://link.springer.com/10.1007/978-981-33-4745-8
10. H. Karunathilake, E. Bakhtavar, G. Chhipi-Shrestha, H. R. Mian, K. Hewage, and R. Sadiq, “Decision making for risk management: A multi-criteria perspective,” vol. 4, pp. 239–287, Jan. 2020, doi: 10.1016/BS.MCPS.2020.02.004.
11. E. Mosqueira-Rey, E. Hernández-Pereira, D. Alonso-Ríos, J. Bobes-Bascarán, and Á. Fernández-Leal, “Human-in-the-loop machine learning: a state of the art,” *Artificial Intelligence Review 2022 56:4*, vol. 56, no. 4, pp. 3005–3054, Aug. 2022, doi: 10.1007/S10462-022-10246-W.
12. K. M. Tay and C. P. Lim, “On the Use of Fuzzy Inference Systems for Assessment and Decision Making Problems,” *Intelligent Systems Reference Library*, vol. 4, pp. 233–246, 2010, doi: 10.1007/978-3-642-13639-9_10.
13. Diego Ongaro and John Ousterhout, *“In search of an understandable consensus algorithm,”* USENIX Annual Technical Conference (USENIX ATC'14), 2014.
14. “HyperLedger Fabric.” Accessed: Apr. 29, 2025. [Online]. Available: https://www.lfdecentralizedtrust.org/projects/fabric
15. “MongoDB.” Accessed: Apr. 29, 2025. [Online]. Available: https://www.mongodb.com/
16. M. Brunelli, “Introduction to the Analytic Hierarchy Process”, *Springer*, 2015.
17. Pandey, V., Komal & Dincer, H. (2023). A review on TOPSIS method and its extensions for different applications with recent development. *Soft Computing 27*, 18011-18039.