

Data sharing in agricultural supply chains: Using semantics to enable sustainable food systems

Christopher Brewster^{a,c,*}, Nikos Kalatzis^b, Barry Nouwt^a, Han Kruiger^a and Jack Verhoosel^a

^a *Data Science Group, TNO, Kampweg 55, Soesterberg, 3769 DE, The Netherlands*

E-mails: Christopher.Brewster@tno.nl, Barry.Nouwt@tno.nl, Han.Kruiger@tno.nl, Jack.Verhoosel@tno.nl

^b *Neuropublic SA, 6 Methonis Str, 18545, Piraeus, Greece*

E-mail: Nikos_Kalatzis@neuropublic.gr

^c *Institute of Data Science, Maastricht University, Maastricht, 6229 EN, The Netherlands*

Editors: Emma Griffiths, Simon Fraser University, Canada; Pascal Hitzler, Kansas State University, USA; Krzysztof Janowicz, University of California Santa Barbara, USA; Monika Solanki, Agrimetrics, UK; Ajay Sharda, Kansas State University, USA; Cogan Shimizu, Wright State University, USA; Evan Wallace, National Institute of Standards and Technology (NIST), USA

Solicited reviews: Emma Griffiths, Simon Fraser University, Canada; Cogan Shimizu, Wright State University, USA; Michael Brückner, Naresuan University, Thailand; one anonymous reviewer

Abstract. The agrifood system faces a great many economic, social and environmental challenges. One of the biggest practical challenges has been to achieve greater data sharing throughout the agrifood system and the supply chain, both to inform other stakeholders about a product and equally to incentivise greater environmental sustainability. In this paper, a data sharing architecture is described built on three principles (a) reuse of existing semantic standards; (b) integration with legacy systems; and (c) a distributed architecture where stakeholders control access to their own data. The system has been developed based on the requirements of commercial users and is designed to allow queries across a federated network of agrifood stakeholders. The Ploutos semantic model is built on an integration of existing ontologies. The Ploutos architecture is built on a discovery directory and interoperability enablers, which use graph query patterns to traverse the network and collect the requisite data to be shared. The system is exemplified in the context of a pilot involving commercial stakeholders in the processed fruit sector. The data sharing approach is highly extensible with considerable potential for capturing sustainability related data.

Keywords: Data sharing, supply chain, agrifood, graph pattern, ontology, Farm Management Systems

1. Introduction and innovation

One of the core challenges faced by the agrifood sector around the world has been the difficulty of sharing data while retaining business confidentiality. This paper presents the technological infrastructure built on semantic web technologies which enable constrained data sharing between multiple agrifood actors. The core innovation lies in the design and building of a framework that enables the querying and translation of data across multiple underlying systems and platforms. To this end we have designed an ontology built on a number of existing agrifood related

*Corresponding author. E-mail: Christopher.Brewster@tno.nl.

standards, and a set of semantic web technology-based modules that enable querying and mapping of data from legacy systems. This has been built in view of the requirements of a frozen fruit supply chain participating as a pilot in the EC funded Ploutos project (<https://ploutos-h2020.eu/>). The proof-of-concept has been designed to be widely reusable (all components are open source and freely available) both in other pilots in the project, and more generally in the agrifood technology sector.

2. Motivation

The agrifood sector is facing a series of immense challenges including economic, social and environmental problems and predicaments. From a financial perspective, agriculture is frequently unprofitable with farmers facing continuous pressures to cut costs, consolidate or find alternative income streams. From a social perspective, there has been a flight from the land for a long time with both push and pull factors that continue across the world today. Recent decades have seen an immense growth in scientific awareness of the environmental impact of agriculture. The agrifood sector is responsible for over 34% of greenhouse gases, uses 70–80% of global freshwater, causes immense biodiversity loss, and is suffering from the annual loss of 27 billion tons of fertile land due to land degradation [10,12,27,44]. The Ploutos project has been designed as a minor contributor to mitigating these impacts. It takes an integrative approach by bringing together innovations in behavioural change and business models, as well data driven technological innovation. The overarching ambition has been to re-balance the value chain and achieve greater sustainability in all three dimensions. Data driven technological innovation largely focuses on ways to achieve greater data sharing across different sets of stakeholders. This paper describes the semantic infrastructure which enables such data sharing among multiple agrifood stakeholders.

Data sharing in this context refers to the design of systems, standards and infrastructure that enables data to be shared between stakeholders both with supply chain relations but also off-chain stakeholders such as certification providers and food safety agencies. The importance of data sharing in the agrifood sector is widely recognised [8,9,15,17,38,46], but for our purpose here it is needed primarily for:

- a. Monitoring purpose e.g. data collection for agricultural certification (GlobalG.A.P., organic etc.) or environmental reporting (of great significance in the new EU Common Agricultural Policy).
- b. Promotional and reputation management purposes e.g. providing information and data to consumers and retailers concerning agricultural practices.
- c. Enabling different value chain relations e.g. providing mechanisms for payments for specific agricultural or environmental practices such as carbon farming or protection of biodiversity.
- d. Tracking and tracing e.g. identifying the location or origin of a specific product, or optimising supply chain flows [3].

There are many other purposes for data sharing in the agrifood sector. Most revolve around increased or better decision support for farmers and other food chain actors. Data sharing for any purpose has usually meant architectures and systems which were centralised either explicitly or implicitly [34]. This is apparent in most solutions which centred around the use of the GS1 standards, and (although seemingly contradictory) most systems built using one form or another of blockchain technologies. The literature (e.g. [30]), but above all our analysis of the real-life agrifood requirements and the barriers and drivers to technology adoption, led us to assume the following principles:

- A. Stakeholders should retain control of their data.
- B. Interoperability must be maintained with conventional software, whether legacy or developed in the future.
- C. Vendor lock-in must be avoided thereby ensuring both flexibility and future proofing of any given part of the overall system

Principle A, control of data, implies that data should remain as much as possible in the hands of the actors concerned (e.g. the farmers), that full control of access should be possible, but also that queries or requests for data should be for specific purposes. Principle B, interoperability with other systems, means there is no expectation that a specific system will be adopted by all stakeholders, only that interoperability with past and future systems should

be built into the design. Principle C, avoiding vendor lock-in, implies that a standards-based, semantic technology driven approach is inevitable, and consequently the use or reuse of widely accepted ontologies.

From a practical perspective data sharing has meant the application of semantic technologies, and to a greater or lesser extent the adoption of the FAIR data principles [45]. As demonstrated below, this has meant the use of ontologies, URIs, and some degree of reasoning. The *Findability* of data is ensured through the use of a registry and a discovery directory, the *Accessibility* of data is ensured by the use of standardised protocols integrated with access control, *Interoperability* of data is ensured by the (re)use of widely used ontologies/vocabularies that are accessible online, and the *Reusability* of data is ensured by using community standards and by ensuring data provenance is a consequence of data ownership and control. Obviously here the target is the relevant agrifood community rather than the wider scientific research community.

The rest of this paper describes the technology developed and deployed for the project and is structured as follows: Section 3 provides an overview of pilot requirements (i.e. needs assessment) and competency questions, followed by a description of the overall architecture in Section 4. Section 5 describes the semantic model developed for the project, followed by a worked example from the frozen food pilot in Section 6. Section 7 describes relevant related work, followed by a discussion of certain limitations in Section 8. This is followed by a conclusion and outline of future work.

3. Pilot requirements and competency questions

Project pilots The scenarios which have driven the technological approach described below have arisen from the pilots of the Ploutos project. This three-year project, running from Oct 2020, aims to integrate innovation in behaviour change and in business modelling together with data driven technologies to help re-balance the agrifood value chain and make the agrifood system more sustainable. The project has involved 33 partners including 22 end-users in 11 pilots, representing a variety of actors in the food system, including farmers, food industry companies, advisors, and ICT providers. The pilots cover a range of agrifood ecosystems, covering arable, horticulture (both open fields and greenhouses), perennials and dairy production. The objective for each pilot is that behaviour change, collaborative business modelling and data driven innovation be integrated and deliver the most environmentally, socially, and economically sustainable solutions possible. Table 1 provides an overview of the Ploutos pilots, the countries involved and the targeted agrifood sector. In each pilot, 3–4 partners are participating, each one providing complementary expertise (e.g. farmers, farmers associations, food processors, retailers and ICT providers).

Requirements analysis process Each pilot has been analysed through a number of interviews with key participants and the use of questionnaires aiming to extract user and technical requirements. From the analysis, it was evident that most of the Ploutos pilots demonstrated the need for more effective mechanisms for “data capturing”, “data driven farm advisory services” and “to establish data flows” between key domain players such as “farmers”, “processors”, “retailers” and “consumers”. Only those pilots which (a) expressed a clear need for sharing and combination of data between different stakeholders, and (b) demonstrated clear requirements for modelling data in a common vocabulary, were chosen so that six out of the eleven pilots have been the initial focus of development. For each of these pilots, an analysis was undertaken of their overall motivations, workflows and high-level architecture. This enabled the identification of the classes which needed to be included in the common semantic model. A set of competency questions was defined as a result, as well as identifying relevant standards (vocabularies, ontologies) that could be reused. Here further details are provided for Pilot 1, as an example, although the full analysis can be found in [40].

Prototypical pilot Pilot 1 is typical in its complexity and requirements. The main pilot participants are located in the Pella region of Northern Greece, and involve two peach farms of about 10 ha in total. A fruit processing factory participates in the pilot in the same area and here the harvested peaches are processed, packaged and sold as a frozen product, mostly to yogurt manufacturers elsewhere in Europe. Farmers have to tackle a number of sustainability problems, as their farms are small, fragmented and in different micro-climate zones, while there is a lack of financial resources for ICT related investments. Current production methods suffer from increasing consumption of inputs and costs with a consequent increased environmental footprint. Fruit processors operating in the area are following

Table 1
Overview of the ploutos pilots

No	Description	Countries Involved	Sector
1	Supporting a frozen fruit value chain with small farmers, to optimise production, reduce environmental footprint and re-use the data for certification and subsidies	Greece	Frozen Fruit
2	Better food-chain contracts for improved durum wheat production	Italy	Arable/Pasta
3	Empowering consumers through crowdsourcing to take back control over their food and create healthy, sustainable, fair trade products	France, Greece, UK, Germany, Belgium	Cross-sectoral
4	Traceability solutions covering the horticulture greenhouses value chain to improve operations, sustainability performance and brand recognition	Spain	Vegetables
5	Smart Farming on rural farms demonstrating its benefit in the wider agrifood community and co-creating new food products and services	Ireland	Livestock, arable/Food tourism
6	Applying soil-passport approach rewarding land owners/users and a precision farming solution to increase soil health and sustainability	Slovenia	Cross-sectoral
7	Supporting wine producers in taking advantage of the changes in labelling regulations and enhancing their sustainability performance	Cyprus	Wine
8	Carbon Farming: compensating farmers for climate friendly soil management	Netherlands	Cross-sectoral/Organic
9	Facilitating the transfer of surplus food from farms to socially disadvantaged groups, by aligning logistics and processes	Serbia, N. Macedonia	Cross-sectoral
10	Increase sustainability in the grapevine sector by introducing payments for ecosystem services provision and parametric insurance to support losses from sustainable approaches	Italy	Fruit
11	Improving the sustainability of Balearics agri-food chains with Smart Farming and by using the collected info to organise agri-tourism activities	Spain	Vegetables/Agri-Tourism

various schemes where better prices are offered for farmers with high quality products. The quality is evaluated by examining the characteristics of the fruits (e.g. taste, shape) but also by checking the cultivation practices followed. It is important for the fruit processors to have demonstrable evidence of the agronomic practices followed (e.g. by means of a “farm book” i.e. diary of actions undertaken) and the exact origin of the produced fruits. Equally the collection of this data is important for a number of certifications processes with regards to good agricultural practices (e.g. GLOBALG.A.P.).

For the purposes of data monitoring at the field level and the optimisation of the farming practices, a commercial Farm Management Information System (FMIS) installed at the two farms was involved in the pilot. This FMIS provides a network of agro-environmental sensors and data collection mechanisms along with digital tools for decision support and digital recording of cultivation practices undertaken. The deployed FMIS is a proprietary system with its own custom data model and with limited capabilities for data sharing. The smart farming solution is supported by a local contractor (agronomist) that also offers technical support and training services to the farmers. The fruit processing factory already uses an Enterprise Resource Planning (ERP) system combined with a Warehouse Management System (WMS) in order to monitor, track and record the processing activities.

The main actors of this pilot are using legacy systems with limited capabilities for interaction (no standardised data formats and APIs) and they are reluctant to proceed with major changes in their operational ICT systems in order to facilitate data sharing. The actors expect to have complete control over their data and who can access it, which means they are unwilling to share complete data sets with third parties.

Given that the raw and processed peaches are mainly exported to European countries, data modelling and data sharing mechanisms need to address not only the needs of the local stakeholders (e.g. at a regional/national level) but to follow current best practices in order to achieve a maximum level of compliance with dominant data sharing practices.

Competency questions Standard competency questions applied to most pilots, including questions, such as “Who owns farm X?”, “What parcels belong to farm X?”, “Are peaches fruit?”, “Give me the GPS coordinates of farm X?”. These questions while trivial are still important to validate the common semantic model. Pilot specific competency questions were identified in each case. For Pilot 1, the following specific questions were proposed:

- It was this detailed bottom-up analysis of pilot requirements, which has led to the adoption of the three principles described in Section 2.

With the digitisation of farm practices and decision support systems (e.g. Farm Management Information Systems), data needs to be shared not just vertically up and down the supply chain but also horizontally e.g. between comparable farmers needing to share data for bench marking or productivity comparisons. The current reality is that data, if exchanged at all, is shared in multiple formats and through various modalities (including manually handing over a USB key). Most existing systems and proposed architectures assume one-to-one interoperability between systems and at most allow for one actor to have access to all data (cf. Fig. 1).

The core technical requirements concern the data collection and data management functionalities necessary for the realisation of the Ploutos pilots. The requirements identified were based on an analysis of the pilots combined with the high level design principles described in Section 2. These technical and software requirements drove the design and implementation of the functional components of the Ploutos data-sharing framework (cf. [20] for more details)

One of the core objectives of the Ploutos data sharing framework is to achieve the controlled and technically sound flow of data between the various information providers and consumers without disturbing the current operations of the underlying systems. In order to achieve this the **Ploutos Interoperability Enabler** (PIE) has been introduced which aims to be generic enough in order to be deployed as a plug and play extendable module on top of the targeted system with minimal customisation efforts. The PIE operates in combination with the **Ploutos Registration and**




Fig. 1. Current practices for data sharing (if it occurs).




Fig. 2. Conceptual view of ploutos data sharing approach.

Discovery Directory (PRDD) (cf. Section 4.4) in order for the various PIEs to announce their existence and their capabilities in order to be discoverable by entities interested in exchanging data. One of the core characteristics of this approach is that data are not shared with a common third party but are exchanged directly among peers. As illustrated in Fig. 2, each data source or sink maintains a specific adaptation of the PIE. The PIE is able to receive queries for data from remote systems, to fetch the appropriate data sets from the underlying system, to translate these data into the common data model (cf. Section 5) and to transfer the data to the remote system. The Ploutos Registration and Discovery Directory acts as a catalogue where the various PIEs are registering their existence i.e. only the accessibility properties (such as hostname and port). Operational properties and capabilities are communicated peer-to-peer. The functionalities of these two core components are elaborated in the following sections.

4.3. Ploutos Interoperability Enabler (PIE)

The Ploutos Interoperability Enabler from a deployment perspective is hosted within the underlying system's (e.g. FMIS, DSS, data collection service) administrative cyber premises as a trusted service while the overall functionality and data sharing is controlled by the administrators of the system. The PIE consists of a set of components which extend the functionality of the underlying systems with specific features. The PIE's main role within the Ploutos architecture is to allow knowledge bases to exchange data in an interoperable manner with other participants of the Ploutos data sharing network. A Knowledge Base can be any service, application or platform that: (1) needs certain knowledge to function, (2) provides certain knowledge that others might need, or (3) both. Examples of Knowledge Bases are: a service that provides a forecast of local temperatures when given a GPS location, an app that gives insight into the supply chain of tomatoes, a platform that manages different sensors on a farm or a database that stores a farmer's farm calendar. A functional component diagram of the PIE along with the potential interactions with external systems is presented in Fig. 3 and analysed in the following paragraphs.

Registration client A PIE will support the underlying "Knowledge Base" by advertising its existence using meta-data descriptions of the respective type of services that it offers (e.g. available information types, data utilisation policies, location, time span of data sets). This registration process occurs in an automated manner. A PIE's capabilities are distributed directly to other PIEs through a peer-to-peer protocol. Additionally, the Registration Client periodically contacts the "Ploutos Registry and Discovery Directory" (PRDD) to provide connection details. The registration process specifies how the exchange of data and knowledge will be realised. For example, the following questions are answered during the registration phase of a PIE:




Fig. 3. Functional component diagram of the Plutos Interoperability Enabler (PIE).

- What knowledge can be requested from me?
- What knowledge will I publish to the network?
- What knowledge will I request from the network?
- To which knowledge will I subscribe?

For example, a temperature sensor might regularly publish temperature measurements to the interested parties, and will respond to requests for the current temperature. A weather station app might subscribe to knowledge about temperature measurements in a field, or request the current temperature. It might also publish current temperature preferences of a user or a specific crop. A FMIS system might subscribe to both the knowledge about temperature preferences of a crop and temperature measurements to be able to optimally recommend interventions such as irrigation.

SPARQL & Messaging API Each PIE maintains a SPARQL & Messaging API ensuring interoperability on a syntactic level. The SPARQL Query Language is a Declarative Query Language (like SQL) for performing Data Manipulation and Data Definition operations on Data represented as a collection of RDF Language sentences/statements. Messaging functionality is achieved using a Developer REST API using the basic graph pattern elements of SPARQL.

Data import REST API This component provides the main mechanisms for the PIE to exchange data sets with the underlying system on which it is deployed. The main gateway of the PIE is a REST API that incorporates the necessary security mechanisms. It is expected that the underlying “Knowledge Base” will consume this REST API to exchange data with the PIE. However, given the expected heterogeneity of the systems that the PIE will interact with additional mechanisms are provided to facilitate the connection with these systems. In case the “Knowledge Base” already provides its own API, the optional component “Knowledge Mapper” is used which acts as an intermediate between the “Data Import REST API” and the API of the underlying system. In case the “Knowledge Base” does not provides any APIs or it is not feasible to consume the PIE REST API an alternative approach is supported by the PIE where a “Data Import Graphical User Interface” allows the manual importing of data.

Knowledge Mapper This enables interoperability at semantic level through a data translation service that will realise the conversion of data streams provided by the hosting system to selected standardised data formats and vice versa. The translation functionality of the “Knowledge Mapper” will be adapted depending on the properties of a selected common data model, however for the needs of the Ploutos project and pilots the translation has been realised according to the classes and properties of the Ploutos Common Semantic Model (PCSM).

Security and access control Within the Ploutos ecosystem ensuring security and access control on data exchange is considered a transversal issue that spans the whole ecosystem. Security related functionality provided by the PIE are based on existing best practices ensuring confidentiality, integrity and availability of data. Moreover, the interaction between PIEs and between a PIE and its Knowledge Base is supported by secure communication protocols, like SSH and HTTPS. In addition, in the current implementation, a role-based access control mechanism is used to enforce access policies on parts of the data provided by a knowledge base. A PIE capability, which defines the subset of data to be provided is defined as the unit of access control. Each knowledge base is owned by an actor or participant that has a certain role in the knowledge network. Access policies are then defined per role as the permission to access a PIEs capability granted to that role. These access policies are maintained by the PRDD and used by the PIEs to enforce them during data exchange phases, e.g. when a remote data query is received by a PIE. The PIE provides the means for enforcing the dictated authentication, access control and data governance policies during data exchange.

Reasoning engine Finally, the PIE contains a reasoning engine that supports two types of reasoning within the Ploutos data sharing framework: (a) reasoning to infer new data and (b) reasoning for orchestration of data exchange. In the first case, rule-based analysis of the collection of existing data/facts allows the inference of new facts and knowledge. As it is described in [6], this is a standard process for the semantic web and a plethora of reasoners are available to be reused. Inferring new data by reasoning is employed in selected cases where the enrichment of the knowledge base is considered useful for the needs of a pilot. The second case of reasoning refers to the orchestration of the process of requesting data from or providing data to knowledge bases and their PIEs in the entire knowledge network. Using this reasoning process, the various PIEs interact with each other to satisfy each other's knowledge needs. These two types of reasoning allows knowledge that resides in separate knowledge bases to be automatically retrieved, transformed and combined, which increases the overall interoperability.

4.4. Ploutos Registration and Discovery Directory

The Ploutos Registration and Discovery Directory (PRDD) is deployed in a server accessible through the Internet. The core objective of this directory service is to allow the registration/discovery of the various PIEs along with their characteristics and to support the orchestration of knowledge discovery. Metadata-like capabilities are not stored in the PRDD, but are communicated directly between PIEs. As already described, data sharing within the Ploutos ecosystem is ensured through the use of an ontology (cf. below). The domain's knowledge model is written in RDF/OWL, which allows reasoning capabilities to be used. The complementary use of PIE and PRDD provides the necessary awareness about the supply and demand of knowledge in the network, allowing the use of reasoning to orchestrate the knowledge supply on-demand. This means that, given a specification of knowledge that is requested, a PIE can infer the appropriate knowledge base to get it. It should be noted, that this approach allows the complementary distributed querying of knowledge bases to serve one query. In addition, given that the PIE is aware of changes in the network, new knowledge bases can be dynamically added to the network. In summary, the use of PRDD and PIE provides the following advantages:

- Knowledge orchestration removes the need to implement compatibility between all pairs of knowledge bases in the network by hand.
- Changes in the knowledge network are handled seamlessly by synchronising information about knowledge interactions.
- Established open-source Semantic Web technologies are leveraged to provide knowledge models and reasoning capabilities.

Figure 4 shows the functional components and the potential interactions of the Ploutos Registration and Discovery Directory.

REST API Communications to and from the PRDD are undertaken through a REST API. The REST API allows PIEs to register their identity and connection information. This metadata is also periodically requested by every PIE to update its internal state according to the latest updates in the Knowledge Network. This metadata is used by a particular PIE to determine which other PIEs should be involved in a data exchange.




Fig. 4. Functional component diagram of the Ploutos Registration and Discovery Directory.

Security and access control This component ensures that only authorised entities are allowed to issue queries and retrieve data on the properties of the registered PIEs. It should be noted that the PRDD only maintains metadata, so even if a malicious reference has been added for a specific PIE, the targeted PIE is the entity that will finally grant access to the requested data. Full details of access control will be described in another paper but employs a combination of standard security mechanisms such as rule-based access control and cryptography for protecting the registry of PIEs.

Knowledge base registry This component handles the registration process of the PIEs. When a new PIE appears on the Ploutos ecosystem it will contact the PRDD in order to announce (1) its existence with Identity information, (2) the way that it can be reached (e.g. IP address, URL, port). Other aspects such as (a) the capabilities that it requires from other PIEs, (b) the capabilities that it offers, and (c) the necessary access policies are communicated peer-to-peer with other PIEs.

5. The Ploutos Conceptual Model

Based on the requirements of the various pilots and their data sharing requirements, a first version of a semantic model with common classes and relations has been developed, the Ploutos Core Semantic Model (PCSM) [40]. In the next subsection, the guiding principles for the modelling and selection of classes and relations in the PCSM will be described. Then, the common classes are briefly listed and the overall design of the PCSM in the form of an ontology is described.

5.1. Guiding principles

For the design of the PCSM, the following guidelines have been applied:

1. The PCSM is based on semantic technologies, such as RDF and OWL, because this is the best way of defining formal semantics and provides the flexibility for modular reuse of existing data models or extending them.
2. The PCSM should be a small, core model that covers the main common classes in the agrifood domain ranging from the farm via the supply chain to the consumer.
3. The classes and relations for the PCSM are selected from the requirements of the pilots. When most of the pilots require a certain class, it is part of the PCSM, e.g. the classes farm, farmer, parcel, and soil that are required by most of the pilots.
4. Existing ontologies that already define the required classes are reused by the PCSM as much as possible. Nonetheless, a specific Ploutos namespace for the PCSM has been defined, namely <https://www.tno.nl/agrifood/ontology/ploutos/common#> prefixed as *ploutos*, that is used to inherit the classes of these existing ontologies.

5. Existing ontologies are only reused when they have a clear formal OWL structure that is publicly available and accessible or downloadable in a .owl, .ttl or .rdf format, for instance at the W3C website or the AgroPortal (<http://agroportal.lirmm.fr>). Consequently, reuse of proprietary ontologies of different projects is avoided.
6. Vocabularies and thesauri/taxonomies that simply define and list a large set of hierarchical terms will not be reused in the PCSM other than using the `rdf:isDefinedBy` property to point to the definition of the class in a vocabulary.
7. Reuse of existing classes and properties in the PCSM is done using the `rdfs:subClassOf` or `owl:equivalentClass` construct for classes and the `rdfs:subPropertyOf` construct for properties.
8. Classes and properties that are required by the PCSM but are not yet part of existing ontologies will be added as classes and properties to the Ploutos namespace.
9. The well-known ontology design pattern called Feature-Observation-Property pattern is used whereby classes are modelled as a `sosa:FeatureOfInterest` or an `sosa:ObservableProperty` to be used in measurements defined as `sosa:Observation`. More details on this pattern can be found in the following subsection.

A consequence of these guiding principles is that existing data models that are not formal ontologies are not explicitly inherited and extended by the PCSM. Unfortunately, extensive data models such as NGSI-LD and vocabularies such as AGROVOC fall into this category, as they are either vocabularies that can still be pointed at using a `rdf:isDefinedBy` or XML-based data models that cannot be reused as existing ontology, because they do not have the RDF-based linked data structure. Finally, the classes and terms of existing ontologies are followed, and thus these terms might be an appropriate fit for the agrifood and supply chain domains, e.g. the term *Observation*.

5.2. Feature-Observation-Property design pattern

One of the guiding principles mentioned above is that of the *Feature, Observation, Property* design pattern. The goal of this principle is to give structure to the different classes and properties that appear in the PCSM. A good structure makes it easier to either *find* a class that you are seeking or *add* a new class that was still missing. The *Feature, Observation, Property* design pattern is based on W3C's SOSA (<http://www.w3.org/ns/sosa/>) ontology and has been applied before in other ontologies. It can be applied in domains where the intention is to use different kinds of sensors to measure properties of things on different levels of abstraction. For example, on the one hand you might want sensors on a drone to measure the drought of a particular parcel on a farm, while on the other hand a farmer might sample individual plants on the parcel and measure their growth. This design pattern can be used for the model to accommodate measurements on both levels of detail (the whole parcel versus the individual plant) and optionally allow them to be correlated at a later stage.

The pattern divides the classes of the ontology into three pillars: features, observations, and properties. Each of the pillars hosts different types of classes that can be related with each other in predefined ways. The Feature pillar consists of feature of interest classes. Within this pillar, classes can be defined as being a member of the class `FeatureOfInterest`, which indicates that the class represents things of which a characteristic is of interest. For example, the class `Soil` and `Plant` are defined within the Feature pillar, because various characteristics of these features are measured for proper management. Within the Feature pillar a large set of classes can arise that are of interest to be monitored.

The Property pillar consists of all kinds of properties of the classes from the Feature pillar. For example, the property `Height` is applicable to the `Tree` class and the `Greenhouse` class and the property `Temperature` is applicable to the air surrounding the farm (i.e., weather), but also to heating pipes in a greenhouse. Other examples of properties are carbon dioxide levels, humidity, colour and radiation. These properties are modelled in the PCSM as classes in order to define the semantics properly and, even more important, to use them in relations to various other classes, using for instance the `hasHeight` or `hasTemperature` relation.

The Observation pillar connects the classes from the Feature pillar with the classes from the Property pillar to model measurements. An observation measures a certain property of a certain feature. In principle, every class in the Feature pillar is a `FeatureOfInterest` that can have an `ObservableProperty` in the Property pillar of which measurements can be done. So, measuring the height of a tree is represented as an `Observation` of the `Height ObservableProperty` (property) with the `Tree FeatureOfInterest` (feature).




Fig. 5. OWL diagram of the Feature-Observation-Property pattern and the reuse of the SOSA ontology for this.

Figure 5 shows the `sos:FeatureOfInterest` class, which enables the entire `sos:Observation` structure to be reused for modelling and instantiating measurements of a `sos:ObservableProperty` of a Ploutos class. Note that this design pattern is a useful way of structuring our model, but not every class must fit these three pillars. For example, treatments of parcels and packaging of meat do not naturally fit into these three pillars and have their own separate Operation classes to represent them.

5.3. PCSM classes and reuse of existing ontologies

The PCSM is a core model in the sense that it only represents the main common classes and properties for the agri-food supply chain domain with a focus on sustainability. These classes and properties are defined in the Ploutos namespace. In total, the PCSM contains around 30 Ploutos classes and about 50 Ploutos properties, data-properties plus object-properties. As much as possible Ploutos classes and properties are being defined reusing classes and properties in existing ontologies. Only when necessary a Ploutos class or property is being newly defined.

To model the common classes of the PCSM, the ontologies listed in Table 2 are reused. Unfortunately, not all data models or ontologies considered were compatible, either because the ontology was not in a standard format, or because it was not dereferenceable. Obviously, there are also other ontologies that might become available soon in the necessary format. For instance, the EPPO Global Database (<https://gd.eppo.int/>) maintained by the EU EPPO organisation contains a lot of plant and pesticide information that can be reused easily once there is an ontology available that contains this information in an RDF and OWL-structure.

The following is a selection of high-level common classes that are required for the pilots and are part of the PCSM. They are selected based on their frequency of occurrences of classes identified in the analysis of all the pilots (cf. Section 3). In general, if a class occurs in three or more pilots, it is included here. Different types of classes to be included are distinguished in the core PCSM. Table 3 provides an overview. For each common class it is also indicated whether the class inherits from an existing class in another ontology either by an `rdfs:subClassOf` or `owl:equivalentClass` relation. Where the table states “none” in the column, the class is a primitive one and newly defined for the PCSM.

5.4. PCSM design

In this section, the design of the PCSM is described focusing on the most common core classes. The reuse of classes of existing ontologies and their properties is also demonstrated.

Table 2
Existing ontologies reused in the PCSM

Prefix	Name	Base URI
ENVO	Environment Ontology	http://purl.obolibrary.org/obo/envo.owl#
org	Organizational Ontology	http://www.w3.org/ns/org.ttl#
FoodOn	FoodOn Ontology	http://purl.obolibrary.org/obo/foodon.owl#
s4agri	SAREF4AGRI	https://saref.etsi.org/saref4agri/
SSN	Semantic Sensor Network	http://www.w3.org/ns/ssn/
SOSA	Sensor Observation Sample Actuator	http://www.w3.org/ns/saso/
OM	Ontology of units of Measure	http://www.ontology-of-units-of-measure.org/resource/om-2/
Weather	BIMERR Weather Ontology	https://bimerr.oit.linkeddata.es/def/weather#

Table 3
The common classes that are defined in the PCSM

Vocabulary Category	Common Class	Inherits from (subClass or equivalentClass)	Specialization
Geographical-related classes	ploutos:Farmer	s4agri:Farmer	
	ploutos:Farm	org:Organization and s4agri:Farm	
	ploutos:Parcel	obo:ENVO00000114 (= agricultural field) and s4agri:Parcel	
	ploutos:Crop	obo:PO0000003 (= whole plant) and s4agri:Crop	
Material-related classes	ploutos:Soil	obo:ENVO00001998 (= soil)	
	ploutos:TreatmentMaterial	None	ploutos:Fertilizer ploutos:Pesticide
Action-related classes	ploutos:Operation	None	ploutos:ProductOperation ploutos:ParcelOperation
	sosa:Observation	None	
Environment-related classes	ploutos:WeatherParameter	weather:WeatherProperty	weather:Temperature weather:Humidity weather:Precipitation weather:Wind weather:Radiation
Certification-related classes	ploutos:Certificate ploutos:CertificateProvider ploutos:CertificationAuditor		

Farmer, farm, parcel and crop For modelling the Ploutos classes farmer, farm, parcel and crop we reused similar classes in the saref4agri ontology and the ENVO ontology. A diagram of how these classes relate can be found in Fig. 6. As can be seen, the `rdfs:subClassOf` property is used to relate the Ploutos classes with the classes of reused existing ontologies. Instances of these Ploutos classes directly inherit the properties of these reused classes. However, where needed, additional properties are defined for Ploutos classes, for instance the property `ploutos:hasArea` for class `ploutos:Parcel`. In addition, the `s4agri:contains` property is reused to model the fact that a farm contains one or more parcels and that a parcel contains zero or more crops. Also important to note is that the ENVO ontology is mostly used to give extra meaning and explanation to a class that is well defined by ENVO. For example, the `ploutos:Parcel` class is an `rdfs:subClassOf` of `obo:ENVO_00000114` which is the class “agricultural field” that has a clear definition in the hierarchical structure of ENVO including synonyms like “cropland” or “grassland”.




Fig. 6. OWL diagram of the ploutos classes farmer, farm, parcel and crop.

Soil In addition to the main common classes farmer, farm, parcel and crop, the next most common class is that of “soil” which is part of the parcel and in which the crop grows. For modelling the Ploutos class `soil`, we reused a similar class in the ENVO ontology, as shown in Fig. 7. As can be seen, the `rdfs:subClassOf` property is used to relate the `ploutos:Soil` class with the ENVO soil class. Thereby, the Ploutos soil class inherits the properties of the ENVO soil class, e.g. the ‘has part’, ‘has quality’ and ‘part of’ properties. These properties show that soil can consist of various other materials, such as ‘clay’, ‘silt’, ‘sand’, ‘minerals’, ‘organic’, and that it may contain chemical entities and nutrients.

Fertiliser and pesticide Apart from water that is irrigated on the soil that is part of the parcel, also other materials and substances need to be applied to either the soil or the crop on the parcel. In order to capture such materials, the class `ploutos:TreatmentMaterials` is introduced. This class contains the class of fertilisers with nutrients that are needed to keep the soil healthy as well as the class of pesticides that are used to protect the crop against diseases and dangerous insects. Figure 8 shows the OWL diagram of these classes and shows that the concentration of a nutrient in a certain fertiliser is modelled using the `om:Percentage` class, that the cost of a fertiliser is expressed in terms of an `om:Measure`. The `om:Measure` class contains a numerical value and a unit, in this case the amount in euros to be paid for a unit of the fertiliser.




Fig. 7. OWL diagram of the Plutos class soil and reuse of ENVO classes.




Fig. 8. OWL diagram of the Plutos class TreatmentMaterial containing fertilisers and pesticides and reusing OM classes.

Operation In order to model action-related classes, the class `plutos:Operation` is introduced. This class can be specialised for various operations on the objects on or at the farm or further down in the agrifood supply chain. The two main properties of a `plutos:Operation` are the `plutos:isOperatedOn` property, that designates to the object on which the operation is performed, and the `plutos:responsibleAgent` property, that indicates the person or organisation that performs the operation. **ProductOperation:** The `plutos:Operation`





Fig. 9. OWL diagram of the Ploutos class operation, ProductOperation and more specific operations.

class has various subclasses of more specific operations, e.g. the `ploutos:ProductOperation`. This class encapsulates all possible operations on a product produced by for instance the farmer, `ploutos:CuttingOperation`, or the food processing factory, `ploutos:PackagingOperation`. See Fig. 9 for an OWL diagram of these classes and their properties. As can be seen, the class `ploutos:ProductOperation` also has an input and output product of the class `ploutos:Product` using the SSN properties `ssn:hasInput` and `ssn:hasOutput`. The `ploutos:Product` class has additionally a property `ploutos:hasOrigin` that indicates the original farm at which the product is produced. **ParcelOperation:** Besides product operations, specific operations on parcels are distinguished using the class `ploutos:ParcelOperation`, see Fig. 10. This class uses the `ploutos:isOperatedOn` property to indicate on which parcel the operation is performed. Furthermore, the class has a datetime on which it is started and a datetime on which it ended. Finally, the `ploutos:hasAppliedAmount` property indicates the amount of a certain treatment material or substance that is used on the parcel in the operation and the `ploutos:isAppliedPer` property indicates the area of the parcel that is treated with this amount of material or substance. Three different parcel operations are defined in the PCSM as subclasses of the `ploutos:ParcelOperation` class.

1. First, the `ploutos:irrigationOperation` class represents all operations for irrigating water on the parcel and does not have any additional properties apart from the generic properties of the `ploutos:parcelOperation`.
2. Second, the `ploutos:FertilizationOperation` class represents all operations for fertilising the soil on the parcel. In addition to the generic properties of the parcel operation, this operation also defines the fertiliser used, `ploutos:usesFertilizer`, the type of fertilization, `ploutos:hasFertilizationType` (e.g. basal fertilization), the application method, `ploutos:hasApplicationMethod` (e.g. broadcasting), and the treated area, `ploutos:hasTreatedArea`, which is an `om:Measure` that indicates the amount of hectares on the parcel that is fertilised.
3. Third, the `ploutos:CropProtectionOperation` class represents all operations for protecting the crop on the parcel. In addition to the generic properties of the parcel operation, this operation also defines the pesticide used, `ploutos:usesPesticide`, the object to which the pesticide is targeted towards,



`ploutos:isTargetedTowards` (e.g. late blight), and the treated area, `ploutos:hasTreatedArea`, which is an `om:Measure` that indicates the amount of hectares on the parcel that is fertilised.

More properties will need to be added in the future if, for example, other types of parcel operations are sufficiently common.

Weather The PCSM also contains weather properties that are required by most of the pilots to monitor the various environmental conditions around the crop on a parcel. To model this, the PCSM has the `ploutos:WeatherParameter` class and reuses the weather properties defined in the BIMERR weather ontology (<https://bimerr.iot.linkeddata.es/>). The `ploutos:WeatherParameter` is therefore made equivalent to the `weather:WeatherProperty` class in the BIMERR weather ontology using the `owl:equivalentClass` property. An advantage of using the equivalent class property here is that the `ploutos:WeatherParameter` automatically inherits all sub-classes of the `weather:WeatherProperty` class containing the actual parameters like Temperature, Humidity, Wind, Storm etc.

In summary, the PCSM is a common semantic model that contains the main classes and their relations for the agrifood domain with a focus on the farmer end. We reused a set of existing ontologies and their classes to represent the Ploutos classes. The PCSM is now being used for the pilots in Ploutos and specifically for the traceability scenario in one of the pilots. In the future, the PCSM will be extended with new classes that are identified as being common for the agrifood supply chain as a whole. The full version of the PCSM is available from a Gitlab repository set-up for the Ploutos project (<https://gitlab.com/Ploutos-project/>).

6. Worked example

The data sharing solution described above has been implemented, deployed and evaluated in the context of Pilot 1 (cf. Section 3). The core objective of this pilot has been to address the existing data related requirements of a frozen fruit value chain. The two major objectives of this pilot have been:

- Support small farmers in optimising the cultivation practices through a data driven smart farming solution (FMIS).
- Increase production transparency and reduce the administrative burden through data sharing mechanisms.

It should be noted that the deployment and use of the proposed data sharing mechanisms took place in real-world operational environment involving the integration of production systems.

6.1. Demonstration steps

In this section a step-by-step analysis of the demonstrator is presented. The demonstration took place during one cultivation season March – September 2021.

Step 1 – Cultivation of peaches with the use of Smart Farming services: In order to improve the cultivation practices and to provide evidence of the applied farming practices the “gaiasense” Smart Farming (FMIS) solution was used (<https://www.neuropublic.gr/en/smart-farming-gaiasense/>) [21]. The gaiasense solution supports the farmers with data-driven advice on fertilisation, pest management and irrigation based on the collection of data from various sources including IoT-enabled agro-environmental stations, satellite derived data, the farmer’s digital calendar (farm book) and on-the-field observations (scouting) of the cultivation. All the aforementioned data are collected in a central cloud repository where they are stored, integrated and processed using data analytic techniques. The outcome of the processing, the proposed recommended actions, is confirmed by experts (i.e. agronomists) in order to generate farming advice focusing on the optimisation of inputs (irrigation, pest management, and fertilisation) tailored to the needs of the targeted parcel’s conditions. The actual farming practices that were applied as a response to the advice is then fed back to the gaiasense system and recorded in the farm book. Figure 11 provides a screenshot of the gaiasense dashboard developed for the needs of Ploutos. The dashboard presents data referring to the selected peach orchard including elements from the farm book (e.g. growth stages, applied pesticides, irrigation), environmental measurements derived by the sensors network and calculated pest infestation risk indexes. Farmers and advisers have access to their respective data collections through web-based dashboards. **Note:** The gaisense FMIS system is not part of the Ploutos architecture but merely an example system that is integrated into the architecture and whose underlying data model is mapped to the Ploutos Common Semantic Model.

The harvesting period for the peaches in this area lasts from the end of the summer until the beginning of autumn. The farmers in the pilot collect the peaches in bins with each bin being assigned a unique identifier in the form of a barcode. The respective records including timestamp, farm-id, farmer-id, bin-id, and each bin’s weight are recorded in the farm calendar and are accessible through the gaiasense FMIS API.¹

The data items that have been identified as relevant for the fruit supply chain and that are available through the FMIS webservice include the following:

- Farm’s details: Country, area, town, farm size, coexisting cultivated crops in the farm, scale/category (e.g. small family farm or industrial agriculture), employees data, Farm’s identifier/name
- Farmer’s details: Name, contact details, Farmer’s identifier
- Farmers’ Cooperative details: Contact details, Number of farmers, financial information, other activities of the cooperative.
- Peach orchard: Parcel’s size, parcel’s polygon, type of cultivated peaches
- Cultivation details: dates of reached phenological stages (first leaves, blossoming, fruiting, harvest date)

¹ See <https://gitlab.com/Ploutos-project> for access details.




Fig. 11. The gaisense web-based dashboard presents data concerning to cultivation practices.

- Farming Operations: Recording of application of pesticides (date, type, dose, target), fertilisers (date, type, dose), irrigation (date, dose)
- Harvesting operation: Date, Farm name (Id), Farmer name (id), Bin id

Step 2 – At the processing factory: The bins with the harvested peaches are brought to the farmers' association warehouse and then transferred to the food processing factory. The processing factory is called Alterra (<https://www.alterra.gr/index.php/en/>) and is located in the same area as the pilot farms and serves the producers for the whole region. Bins received by the processing factory are registered with the factory's Warehouse Management System (WMS) along with the date-time of receipt, peaches type and origin (farmer-id and farm-id). Peaches are then fed to processing lines where they are washed, peeled, sliced, frozen and finally packed into plastic bags that are placed in cardboard boxes of 10–15 kg each. Cardboard boxes produced in the same lot are assigned the same lot-id which is expressed again through a barcode and alphanumeric format. The factory's WMS maintains a registry of the starting and ending times of the process along with the bin-id of the raw peaches and the lot-id of the final product. It should be noted that the specific WMS was not exposing an API but it was feasible to export all necessary records in CSV format. In order to enable connectivity, a simple webservice (REST API) was developed that imports the CSV file and provides access to data through specific calls.

Relevant data items derived from the fruit processing factory are the following:

- Date of delivery of peaches at the factory
- Location of factory (especially with regards to the farm in order to deduce how far the raw peaches have travelled)
- Processing stages
- Date of packaging
- Other profile details about the factory

All these relevant data items are capable of representation in the PCSM whose design is described in Section 5.

Step 3 – Second stage of processing: Pallets with cardboard boxes are exported from Greece to various European countries. The recipients of the processed peaches from the pilot farms are mainly large industrial food processing/production companies that further process frozen peaches in order to produce food products (e.g. fruit yogurt). Food processing companies are eager to get more details on the origins, the applied cultivation practices and the processing states of the purchased frozen peaches. Currently various certification schemes are followed (e.g. GlobalGAP) that aim to ensure the high quality of the products but there is a need for additional data that can provide evidence of the agronomic practices applied. For the needs of the Ploutos pilot, a web-based application for mobile devices has been developed. The industrial consumers of the processed fruits can scan the bar code or type the lot-id and according to the access rights grant can retrieve additional information with regards to the cultivation practices and first-stage processing practices. The query is executed through a distribution of queries to registered systems as described in Section 4. The appropriate information describing the processing steps referring to the fruits in the carton box is retrieved and displayed to the end-user's device.

6.2. Traceability information retrieval scenarios

In the following scenario, Fig. 12 shows the key information providers and the interactions between them. The main information providers in this demo are the “Smart Farming Service” (gaiasense.gr) (operated on behalf of the farmers) and the “Fruit Processing Factory” (alterra.gr). The typical requester of information might be a food processing factory using the frozen peaches. The use case or competency question could be a request from the receiver of the frozen peaches to know when the peaches were harvested or when last sprayed with a pesticide and what type this was. Such information theoretically is available (stored in our example in the gaiasense FMIS as part of the farm calendar) but usually entirely unavailable to third parties. As shown in the sequence diagrams below, there are specific points where data sharing previously was either impossible or involved considerable manual intervention.

The retrieval of information through the agricultural supply chain of this demonstrator is initiated by the issuing of a lot-id based query with the use of the “Ploutos Traceability app”. The overall process is realised in two stages:

- In the first stage, the supply chain is traced backwards to get an overview of the operations that have been performed and the respective identifiers of the participated stakeholders.
- In the second stage, more detailed information from a stakeholder or a specific operation is retrieved and presented upon request by the end-user.




Fig. 12. Conceptual representation of key information providers and interactions of the agri-food data sharing demonstrator.




Fig. 13. Sequence diagram of tracing back into food-chain to retrieve operation information.

Retrieval of overview of processing steps In the context of the realised demonstration, the message sequence diagram in Fig. 13 corresponds to the first stage interactions between the end user's traceability app, the warehouse management system of the fruit processing factory (Alterra) and the Smart Farming (gaiasense) system. The following sequence of steps are identified:

1. The traceability app receives a lot-id from the end-user or the GUI for which an overview of the operations and stakeholders is requested.
2. The traceability app performs a request on its PIE with a graph pattern that expresses an ASK for operations that have directly led to the end-product identified by the lot-id. The following graph pattern is defined and used for the knowledge interactions between the PIEs to get an overview of the operations in the supply chain.

```

?endProduct rdf:type ploutos:Product .
?output ploutos:intermediateProductOf ?endProduct .
?operation rdf:type ?operationType .
?operationType rdfs:subClassOf ploutos:ProductOperation .
?operation ploutos:hasResponsibleAgent ?agent .
?agent s4agri:hasName ?name .
?operation ssn:hasInput ?input .
?operation ssn:hasOutput ?output .
?operation ploutos:hasEndDatetime ?timestamp .
  
```

3. The reasoner in the traceability app PIE uses the PRDD and discovers that the Alterra PIE can provide this information. It then makes a request to the Alterra PIE to ANSWER with its bindings to this graph pattern.
4. The Alterra PIE uses its KM to call the API of the Warehouse Management System to retrieve the list of operations that have led to the lot-id together with the bin_ids that were used as input to the lot-id. This will usually be a packaging operation where fruit in a bin_id is checked and put into a lot-id.
5. Upon receipt of the operations with bin_ids, the reasoner of the Alterra PIE uses the PRDD again to discover that the gaiasense PIE can provide further information on the operations that have led to the products identified by the bin_ids.
6. The gaiasense PIE uses its KM to call the gaiasense API to retrieve, for each bin_id, the operations that have led to this bin_id. This will usually only be a harvesting operation.
7. The gaiasense PIE will send a response back to the Alterra PIE with the resulting list of operations at the farm.
8. The reasoner of the Alterra PIE will combine the list of operations that it received from the gaiasense PIE with the list of operations that it received from the API of the Alterra Warehouse Management System into a complete list of operations.
9. The Alterra PIE will send a response back to the traceability app PIE with the complete list of operations which will then be forwarded to the traceability app.

The result of this sequence of steps is thus a list of “operations” while the actual reply is expressed in JSON format. A reply example follows:

```
[ {
  "endProduct": "<http://alterra.gr/lots/130821A-L4-181011>",
  "operation": "<http://alterra.gr/operations/5488>",
  "input": "<http://alterra.gr/bins/65499>",
  "output": "<http://alterra.gr/lots/130821A-L4-181011>",
  "operationType": "<https://www.tno.nl/agrifood/ontology/ploutos/common#
    PackagingOperation>",
  "agent": "<http://alterra.gr/alterra-org>",
  "agentName": "\"Alterra\"",
  "timestamp": "\"2021-08-24T12:34:56Z\""
},
{
  "endProduct": "<http://alterra.gr/lots/130821A-L4-181011>",
  "operation": "<http://sense-web.neuropublic.gr:8585/operations/1337>",
  "input": "<http://sense-web.neuropublic.gr:8585/crops/23904891289>",
  "output": "<http://alterra.gr/bins/65499>",
  "operationType": "<https://www.tno.nl/agrifood/ontology/ploutos/common#
    HarvestingOperation>",
  "agent": "<http://sense-web.neuropublic.gr:8585/farms/7878>",
  "agentName": "\"Mandalo Farmers Association\"",
  "timestamp": "\"2021-08-22T10:22:00Z\""
}]
```

The reply contains two main items. The first part refers to the ‘creation’ of product <http://alterra.gr/lots/130821A-L4-181011>. The corresponding operation was of type “PackagingOperation” that happened on “2021-08-24T12:34:56Z” where the responsible organisation for this operation was Alterra. This operation belongs to a chain that results in the end product <http://alterra.gr/lots/130821A-L4-181011>, having as inputs the product of this operation the <http://alterra.gr/bins/65499> and output <http://alterra.gr/lots/130821A-L4-181011>. Given that




Fig. 14. Sequence diagram for tracing back to retrieve farm and farming information.

there is a match between the output and the final product this also denotes that it is the end product of this operation. The second part also refers to the creation of product <<http://alterra.gr/lots/130821A-L4-181011>>. The corresponding operation was of type “Harvesting” that happened on “2021-08-22T10:22:00Z” where the entity responsible for this operation was the “Mandalo Farmers Association”. This operation belongs to a chain that results in the end product <<http://alterra.gr/lots/130821A-L4-181011>>, having as input product of this operation the <<http://sense-web.neuropublic.gr:8585/crops/23904891289>> and output <<http://alterra.gr/bins/65499>>.

Retrieving information of a specific processing step – farm and farming practices The sequence diagram in Fig. 14 shows the interactions between the Ploutos traceability application and the Smart Farming service for the retrieval of additional about the farm and the applied farming practices.

Data items of interest that are modelled and retrieved for the needs of the traceability demonstrator include details about the peaches, the farm, the cooperative and harvesting operations. To retrieve this information, the following graph pattern is defined and used for the knowledge interaction between the two PIEs.

```

?operation rdf:type ploutos:HarvestingOperation .
?operation ssn:hasOutput ?bin .
?operation ploutos:isOperatedOn ?parcel .
?parcel s4agri:contains ?crop .
?crop rdf:type ?cropType .
?cropType rdfs:label ?cropName .
?parcel ploutos:hasArea ?area .
?parcel ploutos:hasFirstLeavesDate ?leavesDate .
?parcel ploutos:hasBlossomingDate ?blossomDate .
?parcel ploutos:hasFruitingDate ?fruitingDate .
?parcel s4agri:hasHarvestDate ?harvestDate .
?farm rdf:type ploutos:Farm .
?farm wgs84:location ?location .
  
```

```

?location wgs84:lat ?latitude .
?location wgs84:long ?longitude .
?farm s4agri:hasName ?farmName
?farm s4agri:contains ?parcel .
?farmer s4agri:managesFarm ?farm .
?association rdf:type ploutos:FarmAssociation .
?farmer s4agri:isMemberOf ?association .
?association s4agri:hasName ?associationName.

```

The scope is restricted to *HarvestingOperation* because this is the final operation of the farmer that has as output the harvested peaches put into a *bin_id*. In addition, when the request is made, the *?operation* variable is bound to the value of the harvesting operation_id and the *?bin* variable is bound to one of the *bin_ids* that have led to the original lot-id. The traceability app PIE then enables the traceability app to ASK for knowledge in terms of this graph pattern and the *gaiasense* PIE can ANSWER with their knowledge in terms of this graph pattern. Bearing this in mind, the sequence diagram can be explained as the following sequence of steps:

1. The traceability app receives a request from the end-user or app to get more information about a specific operation with one of the *bin_ids* as output, in this case a harvesting operation at the farm that produced this *bin_id*.
2. The traceability app performs a request on its PIE with a graph pattern that expresses an ASK for the list of crop, parcel, farm and farmer information responsible for this operation.
3. The reasoner in the traceability app PIE uses the PRDD to discover that the *gaiasense* PIE can provide this information and makes a request to the *gaiasense* PIE to ANSWER with its bindings to this graph pattern.
4. The *gaiasense* PIE uses its KM to call the *gaiasense* API to retrieve the requested crop, parcel, farm and farmer information for the harvesting operation of the peaches of the *bin_id*.
5. Upon receipt of this information, the *gaiasense* PIE will send a response back to the traceability app PIE with the information which will then be responded back to the traceability app.

The result of this sequence of steps is thus a set of bindings for the harvesting operation, crop, parcel, farm and farmer information. In JSON format this looks like the following list:

```

[{
  "operation": "<http://sense-web.neuropublic.gr:8585/operations/1337>",
  "bin": "<http://sense-web.neuropublic.gr:8585/bins/65499>",
  "parcel": "<http://sense-web.neuropublic.gr:8585/locations/434>",
  "cropType": "<http://purl.obolibrary.org/obo/NCBITaxon_3760>",
  "cropName": "peach tree",
  "area": "97 hectare",
  "leavesDate": "\"2021-04-22T00:00:00Z\"",
  "blossomDate": "\"2021-05-22T00:00:00Z\"",
  "fruitingDate": "\"2021-06-22T00:00:00Z\"",
  "harvestDate": "\"2021-08-22T00:00:00Z\"",
  "farm": "<http://sense-web.neuropublic.gr:8585/farms/7878>",
  "farmName": "\"Happy Peaches Farm Ltd.\"",
  "latitude": "22.2170236",
  "longitude": "40.7856276"
  "associationName": "\"Mandalo-Proodos\"",
}]

```

In addition, a list of data items of interest to be retrieved for giving insight into the operations on the parcel, soil and plant are the following:

- Farming Operations: Recording of application of pesticides (date, type, dose, target), fertilisers (date, type, dose), irrigation (date, dose)

To retrieve this information, the following graph pattern is defined and used for the knowledge interaction between the two PIEs.

```
?operation rdf:type ?parcelOperationType .
?parcelOperationType rdfs:subClassOf ploutos:ParcelOperation .
?operation ploutos:isOperatedOn ?parcel .
?operation ploutos:hasAppliedAmount ?applied .
?applied om:hasUnit ?appliedUnit .
?applied om:hasNumericalValue ?appliedValue .
?operation ploutos:isAppliedPer ?area .
?area om:hasUnit ?areaUnit .
?area om:hasNumericalValue ?areaValue .
?operation ploutos:hasStartDatetime ?start .
?operation ploutos:hasEndDatetime ?end .
?operation ploutos:responsibleAgent ?agent .
?operation ploutos:usesMaterial ?material .
?material ploutos:hasActiveSubstance ?activeSubstance .
```

In a similar way, a sequence of interaction steps between the PIEs leads to a set of bindings for the parcel operation with more detailed information on fertilizers, pesticides and irrigation. In JSON format this looks like the following list:

```
[{
  "areaUnit": "<http://www.ontology-of-units-of-measure.org/resource/om-2/squareMetre>",
  "area": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/35146dde-efe6-4885-bcfa-28fb67fde45a>",
  "agent": "<http://sense-web.neuropublic.gr:8585/agents/f7b0c242bebc26a65ccaf033b7798298>",
  "parcel": "<http://sense-web.neuropublic.gr:8585/parcels/188451>",
  "parcelOperationType": "<https://www.tno.nl/agrifood/ontology/ploutos/common#IrrigationOperation>",
  "applied": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/471f75bc-a8b8-4f2a-84e3-3dabf4fa5785>",
  "start": "\"2021-06-04T18:00:00\"^^<http://www.w3.org/2001/XMLSchema#dateTime>",
  "appliedUnit": "<http://www.ontology-of-units-of-measure.org/resource/om-2/cubicMetre>",
  "appliedValue": "241.26663906946777",
  "material": "<https://www.tno.nl/agrifood/ontology/ploutos/common#water>",
  "end": "\"2021-06-04T22:00:00\"^^<http://www.w3.org/2001/XMLSchema#dateTime>",
  "activeSubstance": "\"Water\"",
  "areaValue": "4021.1106511577964",
  "operation": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/8cf6ddf2-6e25-44b6-bcc8-c852a3017cad>"
},
```

```

{
  "areaUnit": "<http://www.ontology-of-units-of-measure.org/resource/om-2/squareMetre>",
  "area": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/c41db423-febe-4456-b7a9-075d52182457>",
  "agent": "<http://sense-web.neuropublic.gr:8585/agents/f7b0c242bebc26a65ccaf033b7798298>",
  "parcel": "<http://sense-web.neuropublic.gr:8585/parcels/188451>",
  "parcelOperationType": "<https://www.tno.nl/agrifood/ontology/ploutos/common#CropProtectionOperation>",
  "applied": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/dc75a7d4-14e3-4374-b2eb-9345a47fbb5b>",
  "start": "\"2021-03-12\"^^<http://www.w3.org/2001/XMLSchema#date>",
  "appliedUnit": "<http://www.ontology-of-units-of-measure.org/resource/om-2/gram>",
  "appliedValue": "60.0",
  "material": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/28a58e3cfaa3623888fe64c6da3192b4>",
  "end": "\"2021-03-12\"^^<http://www.w3.org/2001/XMLSchema#date>",
  "activeSubstance": "\"Dithianon\"",
  "areaValue": "4021.1106511577964",
  "operation": "<http://sense-web.neuropublic.gr:8585/.well-known/genid/9bb44bce-3541-4dba-a767-191433374c37>"
}
]

```

A similar approach is taken for retrieving additional data from the food processing factory and for potentially any other stakeholder connected via a PIE.

7. Related work

This section presents related work addressing (a) agrifood related vocabularies and ontologies, (b) data sharing/traceability systems using semantics, (c) querying across federated data using semantics, thereby providing the background and context for the work presented in this paper.

Agrifood related ontologies There is a considerable body of work building ontologies for the food and agriculture domain which has gone hand in hand with the development of Linked Data (and “Linked Open Data”) in the agrifood domain. The major effort here has been AGROVOC developed by the FAO and maintained by a network of institutes around the world [32]. It is currently the most comprehensive multilingual thesaurus and vocabulary for agriculture. AGROVOC has now been partially mapped onto the US National Agricultural Library of the USDA and the CABI thesaurus in the form of the GACS ontology which has mapped and integrated the top 15,000 concepts [4]. Other recent work in this area has also focused on developing ontologies for sharing of research data including the Crop ontology initiative, the Agronomy Ontology (AgrO), and the Plant Trait Ontology (TO) supported by CGIAR [2]. FOODON integrates a number of existing ontologies but its focus seems to be again on research data although its ambition is to provide a mechanism for data integration across the food system. Considerable efforts have been put into extending and integrating the FOODON ontology with various other ontologies extending its utility to areas such as nutrition, and integrating it with the Foodex2 standard from the European Foods Standards Agency (EFSA) [14,16].

As mentioned, most work on ontologies for the agrifood domain has up to now mostly been targeted towards the clear definition of domain concepts and terms in the form of a vocabulary for the annotation of research publications

or research data sets [2]. As a result, there is little or no use yet of ontologies for supporting the actual sharing and exchange of data across the agrifood chain. Only a few papers have been published concerning the use of ontologies for traceability and data analysis in the dairy sector [25,43]. In addition, a few innovation projects have dealt with the use of ontologies in the horticultural supply chain and greenhouse domain [41]. However, the use of semantic standards for information exchange by standardisation organisations, such as GS1 (<https://www.gs1.org/>), ISO or AgGateway (<https://www.aggateway.org/>) is not yet common practice or even beginning to be picked up. More generally under the auspices of the EC funded projects such as ATLAS (<https://www.atlas-h2020.eu/>) and especially DEMETER (<https://h2020-demeter.eu/>), there is support for greater use of semantic standards and technologies where the data and data models are explicitly specified, where URIs are widely used, and where data integration is consequently made far easier [11].

Data sharing architectures The use of semantics for data sharing or data exchange in general has a long history epitomised by the Linked Data initiative and more recently the FAIR data movement [45]. There has been a tendency for this to be largely a focus for researchers rather than commercial applications with the exception of the life sciences. Within the food and agriculture sector, the majority of efforts at data sharing have focused on ontologies for research data and assumed relatively centralised approaches to repositories (e.g. the CGIAR platform for Big Data [2]), or else have been based on XML. Three major data standardisation efforts in the agri-food areas are ICAR (for livestock and dairy production, <https://www.icar.org/>), ISOBUS (for machine-to-machine data sharing, <https://www.aef-online.org/about-us/isobus.html>) and AgGateway (for FMIS applications). These are currently available as XML standards (with codebooks) but are slowly moving towards JSON-LD versions of their standards. [39] make the case for using Linked Data principles and a variety of ontologies so that data can be integrated for farmer decision support. [8] emphasises the importance of semantics for integrating IoT derived data in agriculture, while [42] similarly uses ontologies to ensure data integration for supply chain data. These papers, as do most others, assume centralised architectures. The Linked Pedigree architecture was proposed to enable data sharing across supply chains by formalising the GS1 EPCIS and CBC standards (<https://www.gs1.org/standards/epcis>) as ontologies and enabling SPARQL queries across distributed triple stores, and this work partly inspires the technological approach described in this paper [36,37].

Interoperability and legacy systems The issue of making existing legacy systems interoperable is not new and has been researched as part of the data integration topic [18]. The goal of data integration is focused largely on the need for answering user queries over distributed relational data sources and using SQL as the query language. The data integration landscape has changed with the increase in (types of) data sources and applications that can benefit from it. Other approaches that use semantic technology propose query-rewriting of the overall query to integrate distributed heterogeneous data sources [19]. This is different from our approach in that the data sources are integrated using semantic domain knowledge instead of a user query. Our focus lies on composing web service APIs of the distributed legacy data sources. Applying semantic technologies to web services has been the topic of extensive research. Two topics can be distinguished, although they are often investigated together. On the one hand, research has focused on giving semantic descriptions of the inputs/outputs and pre/post conditions of web services to be able to automatically discover or select relevant services [23]. On the other hand, research focused on how to automatically compose multiple web services together [24]. Despite the many proposals to include semantics into the descriptions of web services [13,26,29], many of today's web service APIs do not describe their formal invocation, pre- and post-conditions, and semantic input/output. The research on semantic web service composition in general has investigated several techniques to achieve this, and the approach described in our paper is similar to rule-based 'planning' [28,33,35]. However, in contrast to these approaches, the approach taken in this paper uses graph patterns of a common semantic model to be able to automatically compose distributed heterogeneous data sources.

In the area of Internet of Things, with the growth of IoT devices efforts have been devoted to overcoming interoperability challenges [1,31]. In this context, the W3C has launched the Web of Things Working Group (<https://www.w3.org/WoT/>) to counter the fragmentation arising from the use of IoT devices and support the use of semantic standards. This work largely builds on the Semantic Sensor Network (SSN) ontology (<https://www.w3.org/TR/vocab-ssn/>). The Open Mobile Alliance and a variety of telecommunication standards bodies have supported the Next Generation Service Interface (NGSI) [5]. The NGSI API and the associated NGSI context model has been

adopted by key organisations in the IoT standardisation efforts. For an example of using this for data translation from multiple sources see [22] which has also influenced the architecture of our interoperability enablers described in Section 4.3.

8. Discussion

The Ploutos data sharing platform enables open and semantically standardised interoperability between existing systems that keep their data at the source. In addition, our approach enables transparent but secure traceability backward up the chain to obtain sustainability information by the retailer or consumer. For the peach example in Section 6, we have created an implementation of the Ploutos data sharing and traceability architecture for the agrifood chain from peach farmer via fruit processor to the food company and consumer. For this purpose, we have developed knowledge mappers and PIEs as well as a traceability app as shown in Fig. 12. A key achievement has been to enable semantic/syntactic interoperability at a high level (e.g. using formal semantic expressions) with existing operational systems without requesting from them to change their mode of operation.

8.1. Feasibility

We used the PCSM to define graph patterns for the knowledge interaction between the traceability app and the farmer's and fruit processor's systems. One of the graph patterns expresses the chain of operations that the peaches go through from harvesting via processing to final packaging and consumption. Another graph pattern is used that expresses more specific sustainability information from the farm, including the extent and use of fertilisers and pesticides. The data derived using this graph pattern at the farmer's system (the FMIS) is made only accessible to the fruit processing company and other intermediaries but not to the consumer (under the current setting of parameters controlling data access).

The implemented PIEs automatically discover and connect to each other based on the knowledge interactions specified. Thereby, the collection of PIEs form a knowledge network that releases the partners in the chain from implementing one-to-one connections with each different partner. The knowledge mappers are used to transform specific information available via an API at the farmer and fruit processor to the PCSM graph patterns. This is the only implementation activity needed to apply the Ploutos data sharing platform to a new use case. A knowledge mapper also has a role-based access control mechanism that permits or denies access to specific graph patterns of the PCSM. The consumer cannot get access to specific parcel operations on the farm, while the fruit processing company can. As a result, the entire traceability scenario shows that it is feasible with limited effort to apply the Ploutos data sharing platform to provide sustainability information to a number of different stakeholders via further PIEs or the Ploutos traceability app.

8.2. Performance

From a scalability perspective, the peach traceability example is fairly small. It only consists of 3 partners with relatively few knowledge interactions. When scaling up to multiple users of (for example) the Ploutos app that request information about a product, the knowledge network should be efficient enough to handle them in parallel. Caching might be one of the approaches here, which is already part of the knowledge mapper implementation. How to increase the level of parallelism and caching in the knowledge network itself is future work.

With respect to response time, a few aspects need to be considered. First, it is important to find the optimum between user actions in the Ploutos app and the number of times knowledge is being collected via the knowledge network. This is fine-tuning work for each application. For example, upon start-up the app can collect all the information that it might show to the user at once, but this might take unacceptably long. An alternative is to only collect information when necessary, e.g. when clicking a button or changing to another tab. Thereby, the total response time is divided over multiple different knowledge interactions.

Second, the number of calls from a Knowledge Mapper to the API of the corresponding knowledge base needs to be limited. In our scenario, we encountered that a lot (pallet) with peaches can consist of peaches coming from 26

bins from the farm. In order to get more information on each of these bins from the farm, 26 calls had to be made to the API. This adds to the total response time of the traceability app querying backwards (upstream) into the supply chain. Splitting up the knowledge interactions might be a way to deal with that, which is an application specific fine-tuning step.

Finally, the number of PIEs (and thus knowledge bases) connected to the knowledge network is a factor that determines the response time. At the moment, a PIE checks upon request of a knowledge interaction which other PIEs exactly match with this knowledge interaction. Subsequently, only matching PIEs will further process the request. As a result, the amount of interactions in the knowledge network is small and the response time depends linearly on the number of PIEs.

9. Conclusions and future work

There has been growing pressure to address the environmental impact of the agrifood sector and with this the recognition that financial, social and environmental dimensions all need to be considered when trying to make the food system less environmentally damaging. The search is on both to incentivise better agronomic practices and to have suitable mechanisms for the monitoring and evaluation of such practices. Technology has been seen both by researchers and policy makers as playing a key role. For example, IoT devices as well as earth observation are seen as providing potential sources of monitoring data. More generally, sensors across the food system in combination with machine learning/AI are seen as providing opportunities not just for monitoring but also prediction and decision support. None of this will be possible without extensive effort in enabling data integration both between devices and systems, and among actors and stakeholders. Heterogeneity of actors makes data sharing a key enabling technology.

This paper has described (1) a bottom-up end-user driven process of requirements collection and technological design in the context of the Ploutos project; (2) the data sharing technology and its architecture, accompanied with an example in a fruit production supply chain; and (3) the Ploutos Core Semantic Model. The requirements collection stage made clear the need for stakeholders to have control over their own data and the uses to which the data is put. Secondly, the importance of integrating any new systems with existing systems is apparent, whether they are legacy systems that a participant already uses or whether they are future systems. Obviously, one of the consequences of designing and making available the PCSM is the hope that some future systems may treat such a data model as a native format or may build in *ex initio* compatible APIs but for this there are no guarantees. The third core requirement concerning the avoidance of vendor lock-in also leads a system designer to choose to use semantic technologies, explicit semantic standards, and as much of the semantic web layer cake of technologies as is useful. The technological components and architecture have been designed to make it possible for a variety of data sharing applications or services to be provided (i) for stakeholders along the supply chain (ii) between actors in a particular stage e.g. among food producers and advisers; and (iii) for off-chain actors such as food safety agencies or similar. An open standards-based approach as proposed here would enable a variety of actors to offer different services to the network while still ensuring the openness enables innovations to occur.

Future work will initially focus on (a) adapting the technical solutions to the needs of the other pilots within the project, including the modular extension of the PCSM for additional agrifood domains; and (b) understanding how the technical solutions provide opportunities for new or different business models, or impose or encourage behavioural changes in order to be successfully adopted, or else, in contrast, provide other new technical requirements. Refining and extending the access control mechanisms by integrating previous work on ontology based access control (OBAC) [7] is also part of planned development, as well developing an easy to use interface for end-users to control their data. It will be important to explore the scalability issues mentioned above including reducing response times to queries, limiting the number of API calls, and more generally designing architectures which can scale up to the numbers of stakeholders found in food systems. Finally, from an adoption perspective, we will further research how and whether the overall approach can be taken up by the wider agrifood technological community.

Acknowledgements

Parts of the work presented here were funded by the European Commission within the H2020 Programme project “PLOUTOS – A Sustainable Innovation Framework to re-balance agri-food value chains” under grant agreement no. 101000594 (<https://ploutos-h2020.eu/>). The authors are grateful to the farmers, agronomists, fruit processors, and other stakeholders involved in the pilots, and we thanks them for their input and feedback.

References

- [1] D. Andročec, M. Novak and D. Oreški, Using semantic web for Internet of things interoperability: A systematic review, *International Journal on Semantic Web and Information Systems (IJSWIS)* **14**(4) (2018), 147–171, ISSN 1552-6283, <https://www.igi-global.com/article/using-semantic-web-for-internet-of-things-interoperability/210657>. doi:10.4018/IJSWIS.2018100108.
- [2] E. Arnaud, M.-A. Laporte, S. Kim, C. Aubert, S. Leonelli, B. Miro, L. Cooper, P. Jaiswal, G. Kruseman, R. Shrestha, P.L. Buttigieg, C.J. Mungall, J. Pietragalla, A. Agbona, J. Muliro, J. Detras, V. Hualla, A. Rathore, R.R. Das, I. Dieng, G. Bauchet, N. Menda, C. Pommier, F. Shaw, D. Lyon, L. Mwanzia, H. Juarez, E. Bonaiuti, B. Chiputwa, O. Obileye, S. Auzoux, E.D. Yeumo, L.A. Mueller, K. Silverstein, A. Lafargue, E. Antezana, M. Devare and B. King, The ontologies community of practice: A CGIAR initiative for big data in agrifood systems, *Patterns* **1**(7) (2020), 2666–3899, <https://www.sciencedirect.com/science/article/pii/S2666389920301392>. doi:10.1016/j.patter.2020.100105.
- [3] J. Astill, R.A. Dara, M. Campbell, J.M. Farber, E.D.G. Fraser, S. Sharif and R.Y. Yada, Transparency in food supply chains: A review of enabling technology solutions, *Trends in Food Science & Technology* **91** (2019), 240–247, ISSN 0924-2244, <http://www.sciencedirect.com/science/article/pii/S0924224418309178>. doi:10.1016/j.tifs.2019.07.024.
- [4] T. Baker, B. Whitehead, R. Musker and J. Keizer, Global agricultural concept space: Lightweight semantics for pragmatic interoperability, *npj Science of Food* **3**(1) (2019), 16, ISSN 2396-8370, <https://www.nature.com/articles/s41538-019-0048-6>. doi:10.1038/s41538-019-0048-6.
- [5] M. Bauer, E. Kovacs, A. Schülke, N. Ito, C. Criminisi, L.-W. Goix and M. Valla, The context API in the OMA next generation service interface, in: *2010 14th International Conference on Intelligence in Next Generation Networks*, 2010, pp. 1–5. doi:10.1109/ICIN.2010.5640931.
- [6] M. Bienvenu, M. Leclère, M.-L. Mugnier and M.-C. Rousset, Reasoning with ontologies, in: *A Guided Tour of Artificial Intelligence Research: Volume I: Knowledge Representation, Reasoning and Learning*, P. Marquis, O. Papini and H. Prade, eds, Springer, Cham, 2020, pp. 185–215. ISBN 978-3-030-06164-7. doi:10.1007/978-3-030-06164-7_6.
- [7] C. Brewster, B. Nouwt, S. Raaijmakers and J. Verhoosel, Ontology-based access control for FAIR data, *Data Intelligence* **2**(1) (2020), 66–77, <http://www.data-intelligence-journal.org/p/36/>. doi:10.1162/dint_a_00029.
- [8] C. Brewster, I. Roussaki, N. Kalatzis, K. Doolin and K. Ellis, IoT in agriculture: Designing a Europe-wide large-scale pilot, *IEEE Communications Magazine* **55**(9) (2017), 26–33, ISSN 0163-6804, <http://ieeexplore.ieee.org/document/8030481/>. doi:10.1109/MCOM.2017.1600528.
- [9] CEMA, Full Deployment of Agricultural Machinery Data-Sharing, Technical Challenges & Solutions, Technical Report, CEMA, 2020.
- [10] M. Crippa, E. Solazzo, D. Guizzardi, F. Monforti-Ferrario, F.N. Tubiello and A. Leip, Food systems are responsible for a third of global anthropogenic GHG emissions, *Nature Food* **2**(3) (2021), 198–209, ISSN 2662-1355, <https://www.nature.com/articles/s43016-021-00225-9>. doi:10.1038/s43016-021-00225-9.
- [11] Demeter_Project, *D2.3 Common Data Models and Semantic Interoperability Mechanisms – Release 2, Deliverable, Deliverable D2.3*, Demeter Project, 2021, <https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=080166e5dc8ed02a&appId=PPGMS>.
- [12] IPBES, S. Diaz, J. Settele, E. Brondizio, H.T. Ngo, M. Gueze, J. Agard and A. Arneth, Summary for Policymakers of the Global Assessment Report on Biodiversity and Ecosystem Services of the Intergovernmental Science-Policy Platform on Biodiversity and Ecosystem Services, Technical Report, IPBES, 2019.
- [13] J. Domingue, D. Roman and M. Stollberg, Web service modeling ontology (WSMO): An ontology for semantic web services, in: *W3C Workshop on Frameworks for Semantics in Web Services*, Innsbruck, Austria, 2005, https://www.w3.org/2005/04/FSWS/Submissions/1/wsmo_position_paper.html.
- [14] D.M. Dooley, E.J. Griffiths, G.S. Gosal, P.L. Buttigieg, R. Hoehndorf, M.C. Lange, L.M. Schriml, F.S.L. Brinkman and W.W.L. Hsiao, FoodOn: A harmonized food ontology to increase global food traceability, quality control and data integration, *npj, Science of Food* **2**(1) (2018), 23, ISSN 2396-8370, <https://www.nature.com/articles/s41538-018-0032-6>. doi:10.1038/s41538-018-0032-6.
- [15] A. Durrant, M. Markovic, D. Matthews, D. May, G. Leontidis and J. Enright, How might technology rise to the challenge of data sharing in agri-food?, *Global Food Security* **28** (2021), 2211–9124, <https://www.sciencedirect.com/science/article/pii/S2211912421000031>. doi:10.1016/j.gfs.2021.100493.
- [16] EFSA, The Food Classification and Description System FoodEx2 (Revision 2), Technical Report, European Food Standards Agency, 2015. doi:10.2903/sp.efsa.2015.EN-804.
- [17] Copa-Cogeca, CEMA, CEETAR, ESA, F. Fertiliser Europe, FEFAC, ECPA, EFFAB and CEJA, EU Code of Conduct on Agricultural Data Sharing by Contractual Agreement, Technical Report, Copa-Cogeca, 2018.

- [18] A.Y. Halevy, A. Rajaraman and J.J. Ordille, Data integration: The teenage years, in: *Proceedings of the 32nd International Conference on Very Large Data Bases*, Seoul, Korea, September 12–15, 2006, U. Dayal, K.-Y. Whang, D.B. Lomet, G. Alonso, G.M. Lohman, M.L. Kersten, S.K. Cha and Y.-K. Kim, eds, ACM, 2006, pp. 9–16, <http://dl.acm.org/citation.cfm?id=1164130>.
- [19] I. Horrocks, M. Giese, E. Kharlamov and A. Waaler, Using semantic technology to tame the data variety challenge, *IEEE Internet Computing* **20**(6) (2016), 62–66, ISSN 1089-7801, <http://ieeexplore.ieee.org/document/7781544/>. doi:10.1109/MIC.2016.121.
- [20] N. Kalatzis, C. Brewster, J. Verhoosel, B. Nouwt, H. Kruiger and A. Georgousis, Initial Agri-Food Data-Sharing Reference Architecture, Deliverable, D4.4, Ploutos Project, 2021, http://ploutos-h2020.eu/wp-content/uploads/2021/09/Ploutos_D4.4-Initial-agri-food-data-sharing-reference-architecture.pdf.
- [21] N. Kalatzis, N. Marianos and F. Chatzipapadopoulos, IoT and data interoperability in agriculture: A case study on the gaiasense smart farming solution, in: *2019 Global IoT Summit (GloTS)*, 2019, pp. 1–6, <https://ieeexplore.ieee.org/abstract/document/8766423>. doi:10.1109/GIOTS.2019.8766423.
- [22] N. Kalatzis, G. Routis, Y. Marinellis, M. Avgeris, I. Roussaki, S. Papavassiliou and M. Anagnostou, Semantic interoperability for IoT platforms in support of decision making: An experiment on early wildfire detection, *Sensors* **19**(3) (2019), 528, <https://www.mdpi.com/1424-8220/19/3/528>. doi:10.3390/s19030528.
- [23] M. Klusch, P. Kapahnke, S. Schulte, F. Lecue and A. Bernstein, Semantic web service search: A brief survey, *KI, Künstliche Intelligenz* **30**(2) (2016), 139–147, ISSN 1610-1987. doi:10.1007/s13218-015-0415-7.
- [24] A.L. Lemos, F. Daniel and B. Benatallah, Web service composition: A survey of techniques and tools, *ACM Computing Surveys* **48**(3) (2015), 33–13341, ISSN 0360-0300. doi:10.1145/2831270.
- [25] L. Magliulo, L. Genovese, V. Peretti and N. Murru, Application of ontologies to traceability in the dairy supply chain, *Agricultural Sciences* **04**(05) (2013), 41–45. ISSN 2156, 8553, 2156–8561, <https://www.scirp.org/journal/paperinformation.aspx?paperid=34577>. doi:10.4236/as.2013.45B008.
- [26] D. Martin, M. Paolucci, S. McIlraith, M. Burstein, D. McDermott, D. McGuinness, B. Parsia, T. Payne, M. Sabou, M. Solanki, N. Srinivasan and K. Sycara, Bringing semantics to web services: The OWL-s approach, in: *Semantic Web Services and Web Process Composition*, J. Cardoso and A. Sheth, eds, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2005, pp. 26–42. ISBN 978-3-540-30581-1. doi:10.1007/978-3-540-30581-1_4.
- [27] J. Mateo-Sagasta, S.M. Zadeh, H. Turrall and M. People, *More Food, Worse Water? A Global Review of Water Pollution from Agriculture, Technical Report, FAO and IWMI*, 2018, <https://www.unwater.org/water-pollution-from-agriculture-a-global-review/>.
- [28] S.A. McIlraith and T.C. Son, Adapting golog for composition of Semantic Web Services, in: *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning, KR'02*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002, pp. 482–496. ISBN 978-1-55860-554-1.
- [29] J. Miller, K. Verma, P. Rajasekaran, A. Sheth, R. Aggarwal and K. Sivashanmugam, WSDL-S: Adding Semantics to WSDL, *Kno. e. sis Publications* (2005), <https://corescholar.libraries.wright.edu/knoesis/199>.
- [30] F. Minnens, N.L. Luijckx and W. Verbeke, Food supply chain stakeholders' perspectives on sharing information to detect and prevent food integrity issues, *Foods* **8**(6) (2019), 225, <https://www.mdpi.com/2304-8158/8/6/225>. doi:10.3390/foods8060225.
- [31] M. Noura, M. Atiquzzaman and M. Gaedke, Interoperability in Internet of things: Taxonomies and open challenges, *Mobile Networks and Applications* **24**(3) (2019), 796–809, ISSN 1572-8153, <https://doi.org/10.1007/s11036-018-1089-9>. doi:10.1007/s11036-018-1089-9.
- [32] S. Rajbhandari and J. Keizer, The AGROVOC concept scheme – a walkthrough, *Journal of Integrative Agriculture* **11**(5) (2012), 694–699, ISSN 2095-3119, <https://www.sciencedirect.com/science/article/pii/S2095311912600586>. doi:10.1016/S2095-3119(12)60058-6.
- [33] D. Redavid, L. Iannone, T.R. Payne and G. Semeraro, OWL-S atomic services composition with SWRL rules, in: *Foundations of Intelligent Systems, 17th International Symposium, ISMIS 2008, Proceedings*, Toronto, Canada, May 20–23, 2008, A. An, S. Matwin, Z.W. Ras and D. Slezak, eds, Lecture Notes in Computer Science, Vol. 4994, Springer, Toronto, Canada, 2008, pp. 605–611, <https://eprints.soton.ac.uk/265132/1/ISMIS08.pdf>.
- [34] H. Scholten, C.N. Verdouw, A. Beulens and J.G.A.J. van der Vorst, Defining and analyzing traceability systems in food supply chains, in: *Advances in Food Traceability Techniques and Technologies*, Elsevier, 2016, pp. 9–33, <http://linkinghub.elsevier.com/retrieve/pii/B9780081003107000028>. ISBN 978-0-08-100310-7. doi:10.1016/B978-0-08-100310-7.00002-8.
- [35] E. Sirin and B. Parsia, Planning for Semantic Web Services, in: *Proceedings of the ISWC 2004 Workshop on Semantic Web Services: Preparing to Meet the World of Business Applications, SWS@ISWC 2004*, Hiroshima, Japan, November 8, 2004, D. Martin, R. Lara and T. Yamaguchi, eds, CEUR Workshop Proceedings, Vols 119, CEUR-WS.org, 2004, <http://ceur-ws.org/Vol-119/paper1.pdf>.
- [36] M. Solanki and C. Brewster, Enhancing visibility in EPCIS governing agri-food supply chains via linked pedigrees, *International Journal on Semantic Web and Information Systems* **10**(3) (2014), 45–73. ISSN 1552-6283, 1552–6291, http://cbrewster.com/papers/Solanki_IJSWIS15.pdf. doi:10.4018/IJSWIS.2014070102.
- [37] M. Solanki and C. Brewster, OntoPedigree: Modelling pedigrees for traceability in supply chains, *Semantic Web* **7**(5) (2016), 483–491, ISSN 1570-0844, <http://www.semantic-web-journal.net/system/files/swj980.pdf>. doi:10.3233/SW-150179.
- [38] K. Spanaki, E. Karafili and S. Despoudi, AI applications of data sharing in agriculture 4.0: A framework for role-based data access control, *International Journal of Information Management* **59** (2021), 102350, ISSN 0268-4012, <https://www.sciencedirect.com/science/article/pii/S0268401221000438>. doi:10.1016/j.ijinfomgt.2021.102350.
- [39] F.K. van Evert, S. Fountas, D. Jakovetic, V. Crnojevic, I. Travlos and C. Kempenaar, Big data for weed control and crop protection, *Weed Research* **57**(4) (2017), 218–233, ISSN 1365-3180. doi:10.1111/wre.12255.
- [40] J. Verhoosel, C. Brewster, H. Kruiger and B. Nouwt, Data Interoperability for the Agri-Food Sector, Deliverable, D4.2, Ploutos Project, 2021, http://ploutos-h2020.eu/wp-content/uploads/2021/09/Ploutos_D4.2-Data-interoperability-for-the-agri-food-sector_v1.0_2021-06-11.pdf.

- [41] J. Verhoosel, M. Van Bekkum and T. Verwaart, Semantic interoperability for data analysis in the food supply chain, *International Journal on Food System Dynamics* **9** (2018), 101–111, <http://centmapress.ilb.uni-bonn.de/ojs/index.php/fsd/article/view/917>. doi:10.18461/IJFSD.V9I1.917.
- [42] J. Verhoosel, M. van Bekkum and T. Verwaart, Semantic interoperability for data analysis in the food supply chain, *International Journal on Food System Dynamics* **9**(1) (2018), 101–111, ISSN 1869-6945, <http://centmapress.ilb.uni-bonn.de/ojs/index.php/fsd/article/view/917>. doi:10.18461/ijfsd.v9i1.917.
- [43] J.P.C. Verhoosel and J. Spek, Applying ontologies in the dairy farming domain for big data analysis, in: *Joint Proceedings of the 3rd Stream Reasoning (SR 2016) and the 1st Semantic Web Technologies for the Internet of Things (SWIT 2016) Workshops Co-Located with 15th International Semantic Web Conference (ISWC 2016)*, Kobe, Japan, October 17–18, 2016, D. Dell’Aglia, E.D. Valle, T. Eiter, M. Krötzsch, M. Maleshkova, R. Verborgh, F.M. Facca and M. Mrissa, eds, CEUR Workshop Proceedings, Vol. 1783, CEUR-WS.org, 2016, pp. 91–100, <https://ceur-ws.org/Vol-1783/paper-09.pdf>.
- [44] S. Wang, Agrobiodiversity and agroecosystem stability, in: *Introduction to Agroecology*, C.D. Caldwell and S. Wang, eds, Springer, Singapore, 2020, pp. 137–154. ISBN 9789811588365. doi:10.1007/978-981-15-8836-5_10.
- [45] M.D. Wilkinson, M. Dumontier, I.J.J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L.B. da Silva Santos, P.E. Bourne, J. Bouwman, A.J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C.T. Evelo, R. Finkers, A. Gonzalez-Beltran, A.J.G. Gray, P. Groth, C. Goble, J.S. Grethe, J. Heringa, P.A.C. ’t Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S.J. Lusher, M.E. Martone, A. Mons, A.L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S.-A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M.A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao and B. Mons, The FAIR guiding principles for scientific data management and stewardship, *Scientific data* **3** (2016), 2052–4463, <http://dx.doi.org/10.1038/sdata.2016.18>. doi:10.1038/sdata.2016.18.
- [46] M. Wysel, D. Baker and W. Billingsley, Data sharing platforms: How value is created from agricultural data, *Agricultural Systems* **193** (2021), 103241, ISSN 0308-521X, <https://www.sciencedirect.com/science/article/pii/S0308521X21001943>. doi:10.1016/j.agry.2021.103241.