

Semantically enriched crop type classification and Linked Earth Observation Data to support the Common Agricultural Policy monitoring

Maria Rousi, Vasileios Sitokonstantinou, Georgios Meditskos, Ioannis Papoutsis, Ilias Gialampoukidis, Alkiviadis Koukos, Vassilia Karathanassi, Thanassis Drivas, Stefanos Vrochidis, Charalampos Kontoes, and Ioannis Kompatsiaris, *Senior Member, IEEE*,

Abstract—During the last decades, massive amounts of satellite images are becoming available that can be enriched with semantic annotations for the creation of value-added Earth Observation products. One challenge is to extract knowledge from the raw satellite data in an automated way and to effectively manage the extracted information in a semantic way, to allow fast and accurate decisions of spatio-temporal nature in a real operational scenario. In this work we present a framework that combines supervised learning for crop type classification on satellite imagery time-series with Semantic Web and Linked Data technologies to assist in the implementation of rule sets by the European Common Agricultural Policy (CAP). The framework collects georeferenced data that are available online and satellite images from the Sentinel-2 mission. We analyze image time-series that cover the entire cultivation period and link each parcel with a specific crop. On top of that, we introduce a semantic layer to facilitate a knowledge-driven management of the available information, capitalizing on ontologies for knowledge representation and semantic rules, to identify possible farmers non-compliance according to the Greening 1 (Crop Diversification) and SMR 1 rule (protection of waters against pollution caused by nitrates) rules of the CAP. Experiments show the effectiveness of the proposed integrated approach in three different scenarios for crop type monitoring and consistency-checking for non-compliance to the CAP rules; the smart sampling of on-the-spot-checks; the automatic detection of CAP's Greening 1 rule; and the automatic detection of susceptible parcels according to the CAP's SMR 1 rule.

Index Terms—semantic enrichment, linking EO data and Web content, crop type classification, EU CAP non-compliance checking, data fusion for decision-making.

I. INTRODUCTION

IN recent years, a massive quantity of georeferenced data is generated from many different sources like human activity and Earth Observation (EO), in-situ sensors, satellite missions

(e.g. Copernicus) and mobile phones. The semantic enrichment and linking of these free and open data of this scale, frequency, and quality constitute a fundamental challenge for interoperability and automation in decision-making. EO data become useful only when analyzed together with other sources of data (e.g. geospatial data, in-situ data) and turned into actionable information and knowledge for decision making. In this context, linked data¹ is a data paradigm that studies how one can make Resource Description Framework (RDF) [1], [2] data available on the web and interconnect it with other data with the aim to increase its value. In the last few years, linked geospatial data has received attention as researchers have started tapping the wealth of geospatial information available on the web using semantic web technologies [3], [4]. Nevertheless, there are only a handful of applications that showcase the semantic integration of linked EO and non-EO products. The scalability to accommodate big linked EO data also remains an open issue [5].

One of the domains that is already heavily dependent on the effective and efficient knowledge extraction from EO data is the control of the Common Agricultural Policy (CAP) [6]. The European Union (EU), through the CAP, aims at increasing the European agricultural productivity under sustainable practices, while at the same time making sure that the farmers maintain a decent standard of living². It is the EU's aim to reinforce the competitiveness of European agriculture, whilst maintaining and strengthening its sustainability. This manifests as a major priority, with CAP's annual budget amounting to approximately 59 billion Euros. The Integrated Administration and Control System (IACS) of the CAP, consumes the majority of its annual budget. The IACS functions as the management system for the CAP payments; and is implemented by the national paying agency of each EU Member State (MS)³. The CAP legal framework is transitioning to its new form, the post-2020 CAP reform, which aims to modernize and simplify the current operating model⁴. Based on the post-2020 CAP ambitions and towards the direction of the so-called monitoring approach for the implementation of IACS, Earth

M. Rousi, G. Meditskos, I. Gialampoukidis, S. Vrochidis, and I. Kompatsiaris are with the Centre for Research and Technology Hellas (CERTH), Information Technologies Institute, Thessaloniki, 57001, Greece e-mail: mariarousi@iti.gr, gmeditsk@iti.gr, heliasgj@iti.gr, stefanos@iti.gr, ikom@iti.gr.

V. Sitokonstantinou, I. Papoutsis, A. Koukos T. Drivas, C. Kontoes are with the National Observatory of Athens (NOA). e-mail: vsito@noa.gr, ipapoutsis@noa.gr, akoukos@noa.gr, tdrivas@noa.gr, kontoes@noa.gr.

V. Karathanassi is with the Laboratory of Remote Sensing, National Technical University of Athens (NTUA), 9 Heroon Polytechniou Str., Zographos, 15790, Athens, Greece e-mail: karathan@survey.ntua.gr (V. Sitokonstantinou is also affiliated with the Laboratory of Remote Sensing, National Technical University of Athens)

¹<https://www.w3.org/standards/semanticweb/data>

²<http://esa-sen4cap.org/>

³https://ec.europa.eu/info/food-farming-fisheries/key-policies/common-agricultural-policy/financing-cap/controls-and-transparency/managing-payments_en

⁴<https://www.consilium.europa.eu/el/policies/cap-future-2020/>

Observation has been identified as a key enabler.

Multiple EC funded projects have employed Copernicus data, using advanced ICT and Artificial Intelligence (AI) technologies, to address the monitoring of the CAP. The RECAP project⁵ has been one of the first to develop Copernicus based machine learning pipelines to assist the Paying Agencies in reducing the costs and increasing the efficiency of the control of CAP's Cross-Compliance. Additionally, the Sen4CAP project⁶, building on the legacy of RECAP, has focused on reducing the costs of IACS towards the post-2020 CAP objectives; exploring the applicability of an evidence-based monitoring approach. To the best of our knowledge, none of the existing approaches are able to support the following CAP scenarios to deliver true business value. The selected scenarios require significant human resources, with On-The-Spot Checks (OTSCs) on only a small sample of the farmers' applications.

a) Smart Sampling of on-the-spot-checks (OTSCs):

Farmers declare the cultivated crop type for their arable land each year, around the months of May and June. Paying agencies are responsible for validating the declarations to then grant the requested subsidy to the farmer. The MSs randomly sample and inspect 1-5% of the total number of declarations. It is necessary to automatically monitor the farmers' declarations, using and linking additional and online available data, in order to create an targeted, instead of random, sample set that requires OTSC.

b) Automatic detection of CAP's Greening 1 rule: This rule seeks to improve biodiversity and reduce soil erosion by imposing limits in the size and number of the different cultivations in a farm. Specifically, the farmers that own 10 to 30 ha of arable farm should grow at least two different crop types, while farmers that own more than 30 ha should grow at least three different crop types. In the first case, the main crop should not cover more than 75% of the land, while for the latter case the two main crops should additionally not exceed 95% of the total land⁷. This rule requires accurate and semantically enriched geospatial data, so as to both detect correctly the crop types and perform semantic reasoning to infer the consistency between the farmers' declaration and the Greening 1 rule.

c) Automatic detection of susceptible parcels according to CAP's SMR 1 rule: This particular CAP requirement expects from the farmers, among other things, to perform a risk assessment on the susceptibility of their parcel to contribute nitrate-rich soil to nearby surface waters. The farmer should account for the slope of land, the ground cover, the proximity to surface water, weather conditions, soil type and conditions and the presence of land drains⁸. The SMR 1 requirement defines buffer zones, which shall be respected in terms of fertilizer spreading. Specifically, i) manufactured fertilizer spreading should be at least 2 m from surface water

and ii) organic manure spreading should be at least 10 m from surface water⁹. Therefore, measuring the proximity of the parcel boundaries to the nearest surface water is of great significance both for the inspections of the paying agency, but also for the farmer who wishes to better comply with the requirement. To that end, the semantic data fusion of Sentinel images and other linked open geospatial data, would lead to a more efficient monitoring of SMR 1.

In this work we adopt and extend the linked open EO data life cycle paradigm [3], and provide for the first time an end-to-end implementation to address operational needs of the CAP. Motivated by the above mentioned use cases and their inherent requirements to effectively and intelligently combine and interlink various geospatial data sources (e.g. Land Parcel Identification System (LPIS), hydrographic network, Natura2000 zones etc.), we instantiate the linked open EO data life cycle in the domain of CAP. To this end, we propose a hybrid data- and knowledge-driven framework, developing concrete CAP-related scenarios and demonstrate automatic pipelines for satellite data processing, content extraction, semantic annotation and transformation to RDF, interlinking layer, validation and querying of Linked Open spatio-temporal data.

Our contributions are summarized as follows:

- We propose the use of semantic reasoning in the context of CAP monitoring, to check compliance with Greening-1 requirements, taking into account satellite derived products and ancillary geospatial data.
- We demonstrate the use of spatial relationships in LOD (GeoSPARQL) towards assessing vulnerable parcels according to CAP SMR-1 specifications.
- We propose the Smart Sampling Scheme, i.e. the use of spatiotemporal queries to define a new, educated sampling of parcels that need to be checked for compliance with CAP rules with in-field visits.
- We evaluate our framework under the light of national scale application, in line with post-2020 CAP monitoring needs. Therefore, we discuss scalability implications for both the knowledge extraction from satellite imagery module and the semantic reasoning framework.

The rest of the paper is structured as follows. Section II presents related technologies considering both the Earth Observation based CAP monitoring and semantic web technologies. The technologies are grouped under a common framework, aiming to cover the whole life-cycle of the linked EO data for the control of the CAP. Section III describes our proposed semantically enriched crop type classification model for checking the compliance of farmers' declarations to the CAP regulations. Section IV presents experiments and results regarding our proposed methodology, along with the area of interest, the considered CAP scenarios, their implementations and results regarding the effectiveness and efficiency of our pipeline. Section V concludes our paper.

⁵<https://www.recap-h2020.eu/>

⁶<http://esa-sen4cap.org/>

⁷https://ec.europa.eu/info/food-farming-fisheries/key-policies/common-agricultural-policy/income-support/greening_en

⁸GCCE 2017 v1.0, "The guide to cross compliance in England 2017", produced by the Department for Environment, Food and Rural Affairs

⁹<https://www.gov.uk/guidance/using-nitrogen-fertilisers-in-nitrate-vulnerable-zones>

II. LIFE-CYCLE OF LINKED EO DATA FOR THE CONTROL OF THE CAP

Existing works for the monitoring of CAP mainly focus on knowledge extraction technologies, while the semantic technologies cover more generic agricultural needs. However, our main objective is to cover the complete life cycle of linked open EO data paradigm, as originally discussed in [3], following a multi-disciplinary approach. The stages of the life-cycle of linked EO data for the control of the CAP are presented in Figure 1. The first step covers the content extraction machine learning methodologies so as to get new information layers out of the large streams of raw satellite data. The second step involves the standardized data representation and ontological modeling for semantic annotation. The semantic annotation is based on semantic Web technologies which are being adapted and developed under the EO and Agricultural domains. The next step of the life-cycle regards the transformation of the extracted content into RDF, allowing the population of the knowledge base (triplestore) to perform semantic queries that offer a better knowledge of the data (Storage/Querying). Applying useful interconnections in the semantic data using external datasets can additionally enrich the content and extract hidden knowledge (Interlinking).

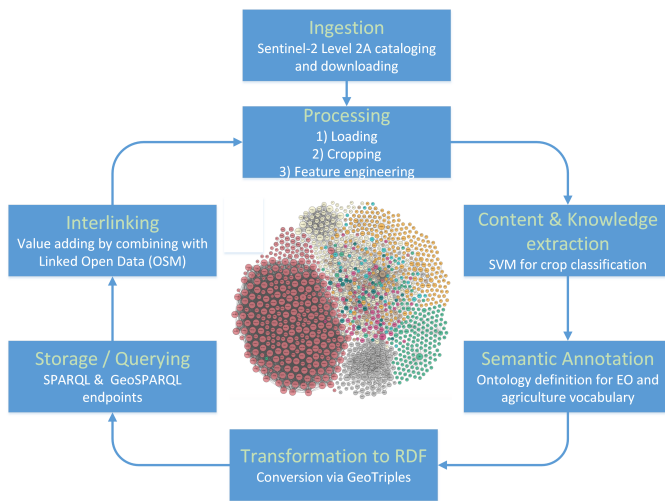


Fig. 1: The main stages of the complete life-cycle of linked open EO data for the monitoring of the CAP

A. Content and Knowledge Extraction: Crop Classification

Over the last decades there have been multiple studies that have utilized EO data to extract high level thematic knowledge for the agricultural land. Recently and since the introduction of the Sentinel missions, there have been a plethora of scientific publications that have exploited either Sentinel-1 or Sentinel-2 imagery, or in certain cases both, to classify crop types. The high spatial and temporal resolution of the Sentinel missions makes them ideal for constructing dense image time-series of high quality that capture all the phenological stages of the different crop types and thus allowing for their accurate discrimination.

The state-of-the-art in EO-based and specifically the Sentinel-based crop classification has advanced significantly over the last years, with the majority of publications reaching optimal accuracy levels ($>85\%$) for multi-class problems. The relative differences in the published approaches are based on the nature and level of specificity of the investigated crop classes, the computational complexity restrictions, the scale of application and the ground truth information that is available for training and validation.

In order to reduce the computational complexity of crop classification and develop scalable solutions, multiple studies have followed Object-based Image Analysis (OBIA) approaches. For instance, in [7] the authors have segmented their image stack into objects using spectral segmentation techniques on Very High Resolution imagery, while in [8] the authors made use of the LPIS to partition their feature space into parcel objects.

Synthetic Aperture Radar (SAR) and optical imagery, retrieved from Sentinel-1 and Sentinel-2 missions respectively, have been used either individually or combined. In [9] a combination of both Sentinel-1 and Sentinel-2 is used in order to create very dense time-series, thus alleviating the cloud coverage limitations. In [10], the authors employ solely Sentinel-1 data, suggesting a weather independent crop classification scheme for the monitoring of the CAP, hence accounting for northern European countries that suffer from year round cloud coverage. Other studies focus on generating multiple diverse features from Sentinel imagery, beyond the most common spectral bands and vegetation indices. In [11] and [12], the authors create deep feature spaces, additionally including variations of vegetation indices, texture and phenology parameters. Such methods are shown to be particularly useful in classifying spectrally heterogeneous crop classes, i.e. vegetables.

With respect to the classification methods employed, both supervised and semi-supervised learning approaches can be found in literature. In [13], for example, the authors have combined a hierarchical correlation clustering with an artificial neural network. The vast majority of studies, however, make use of supervised learners, such as Support Vector Machines (SVM) and Random Forest (RF) ([8], [7], [11], [12], [9], [14], [15]). Their effectiveness stems from their ability to accurately describe the nonlinear relationships between crops' physical condition and their spectral characteristics, while being particularly insensitive to noise and overfitting. Finally, there are important studies that have used Convolutional Neural Networks (CNN) or Recurrent Neural Networks (RNN) or a combination of both [16], which allow the learning of time and space correlation over the Sentinel time-series, thus reducing manual feature engineering.

In [8], the authors have developed a scalable crop identification scheme, employing a 2nd order polynomial SVM on a time-series of Sentinel-2 data. The authors of [8] have additionally performed an extensive comparison between SVM and RF, them being the most widely used classifiers for crop mapping problems. The results showcased the superiority of SVM over RF for the classification of multiple and spectrally similar classes. This conclusion is additionally supported by [17].

This study builds on the methods and results that were published in [8], following the state-of-the-art in crop classification as described earlier. The crop classification method of this study, however, has been applied to three different areas of interest, of diverse characteristics, thus proving its transferability. Finally, we perform multiple crop classifications, starting from very early in the year; therefore with truncated feature spaces. Nonetheless, the results, even early in the year, are satisfactory for the purposes of smart sampling the CAP OTSCs.

B. Semantic Web Technologies

Web Ontology Language (OWL) [18] is an ontology language that provides classes, properties and individuals under the semantic web aspect. Ontologies offer the taxonomy of semantic objects and the relationship between them. RDF [19] [20] is the W3C recommendation standard that offers data representation under subject–predicate–object standard, which is known as triples. Each subject is a resource and each object can be either a resource, a value or an empty node. Predicates or properties express the relationship between a specific subject and object. Data expressed in this format are saved into RDF triplestores, named Knowledge Graphs.

SPARQL Protocol and RDF Query Language (SPARQL) [21] is the most popular querying language for the retrieval and manipulation of data in RDF format. SPARQL offers a wide range of query forms and operators to access and retrieve the data. stSPARQL [22] [23] is a SPARQL version that applies semantic queries into data in stRDF format. Such formats offer representation and querying of thematic and spatial data which contain a temporal dimension. GeoSPARQL [23] focuses more on geospatial data querying, by providing a wide list of functions to support semantic queries execution over geometry and feature objects. Topological relationships are also taken into account.

Reasoning [24] is the procedure of inferring logical consequences based on asserted facts or axioms. In RDF graphs, reasoning takes advantage of data triples using in many cases different data sources to specify the rules that can lead to useful knowledge extraction. Reasoning with rules is usually based on Description Logics (DLs). Description Logics [25] constitute a family of logic-based representation formalisms and are usually used to represent well-structured knowledge over the application domain. Its name comes from a combination of descriptions, which are the expressions namely predicates, and the fact that they support logic-based semantics. DLs are strongly associated with structuring ontology languages such as OWL, but are also widely used in application domain.

The vision of linked data is associated with the transformation of data into RDF formats. Data in this format can be published on the web and linked with other existing data that come from different sources [20]. Linked data are easily accessible using semantic queries. The main advantage of semantics is that they have the means to create intelligent interconnections over objects that come from heterogeneous sources as they support better information management, complexity limitation

and useful inferences extraction [26]. In this work we provide a list of functions and semantic queries, using semantic web technologies, to support three CAP-related scenarios in real operational problems, that require content extraction and semantic linking of data for compliance checking.

C. Semantic Annotation under the Earth Observation and Agriculture domains

Building appropriate ontologies to describe the different aspects of earth observation and agriculture are presented in this section. Earth observation ontologies focus more on the environmental monitoring domain.

The ontology presented in [27] deals with hydrological monitoring issues and captures the main components of hydrological monitoring which are the events, the sensors and the observations. Sensors and observations are divided into many subcategories such as physical and meteorological, while events are associated with any hydrological cycle change. Modular Environmental Monitoring Ontology (MEMOn) [28] extends the abovementioned ontology as, except from sensor and observation data, it provides a structure to model a plethora of different aspects that are identified on an emergency situation under the environmental monitoring domain. The ontology provides the structures to represent environmental features (procedure and material), physical conditions (disaster) and spatiotemporal information (geolocation and time).

An agricultural ontology representation method is described in [29]. The suggested model contains information such as cultivation and processing practices, storage, pests control, genetic attributes etc. OntoCrop ontology [30] offers knowledge representation for common cultivation practices, pests control and in general the crops physiology. Each plant is characterized by properties such as name, growth stage, type of infection, infected part, information about disorders. Agriculture ontology for the purpose of agriculture internet of things (AgOnt) [31] presents a more product-oriented view of agricultural products containing information related to the product, the seeding procedures, the physical conditions, the phase, the location and temporal dimensions. The main purpose of this ontology is supporting healthy food management. The ontology presented in [32] offers a uniform representation of text classification and concepts extraction results. The ontology matches specific concepts into ontology classes which include many different types of products such as agricultural, planting, livestock, fishery and agricultural material.

Most works that have been previously mentioned describe either earth observation or agricultural data. What is actually missing is a combination of both describing information from earth observation with information from the agricultural domain. In this work we reuse and extend the ontology found in [8] that describes the relationship among crop types, families and season, and create a combination with GeoSPARQL vocabulary that represents geospatial-related data such as points in polygon geometry.

D. Transformation into RDF

Another widely investigated issue is combining semantics with earth observation data to discover hidden knowledge. This section describes some frameworks that deal with data transformation into semantic format, integration and searching.

Intelligent Interactive Image Knowledge Retrieval (I^3KR) [33] is a framework that utilizes EO data archives and applies image segmentation (PCA kernel approach) and classification techniques (SVM learning method). From the semantic aspect, the system achieves high-level query processing into context information from distributed data archives. Domain-specific ontologies provide the appropriate structures to integrate heterogeneous data sources in order to support complex semantic queries. A hybrid ontology approach has been used to integrate data coming from different ontologies. Semantic restrictions have been applied using DL reasoning to determine the conditions under which an instance will belong to a class.

GeoTriples [23] is a tool that deals with the geospatial data transformation into semantic RDF format. The system gets as an input a file in various formats and creates a mapping which is based on GeoSPARQL vocabulary, using RML and R2RML rules (mapping generator). Users have the opportunity to define the rules if needed. Initial data are transformed into RDF graph format using the RML rules defined in previous phase (mapping processor). Various RDF syntax formats are supported. Querying is also available in a relational database using R2RML mapping (stSPARQL/GeoSPARQL evaluator).

In this study, we use the GeoTriples tool as a basis in order to transform shapefile data into RDF format under the GeoSPARQL standard for semantic representation.

E. Storage and querying

In this work, we choose to handle the three different CAP scenarios using semantic technologies. The problem could have been solved using relational databases, though this selection would be accompanied with an inflexible data schema and higher execution times [34]. Additionally, the relationships between the entities handled in this work are quite complex to be represented using SQL keys¹⁰. Information is coming from three layers and, with the usage of semantics, are combined in the most effective way, while OWL 2 RL rules are used to enrich the data [35].

RDF triplestores are semantic databases that offer data saving in semantic graph format. Strabon [36] is a geospatial-oriented RDF triplestore that offers a broad amount of querying functions over georeferenced information, supporting both stSPARQL and GeoSPARQL. GraphDB¹¹ [37] is also a popular triplestore that supports saving and querying over georeferenced and non-georeferenced semantic data, supporting native OWL 2 reasoning. It is considered as one of the best triplestores available in terms of storage, supported functionalities, performance and execution time [38]. Other RDF triple stores that provide geospatial support include

RDF4J¹², Virtuoso¹³ [39], OntopSpatial¹⁴ [40], Oracle spatial and Graph¹⁵, AllegroGraph¹⁶, Stardog¹⁷, uSeekM¹⁸ and Parliament¹⁹.

The storage and query capabilities of our framework capitalise on an existing RDF triple store, on top of which SPARQL and GeoSPARQL standards are used to form the queries that support the rules of agriculture policies. The current implementation uses the GraphDB semantic graph database, taking full advantage of the provided dashboard to explore and manage the RDF repositories. It also supports different reasoning profiles, such as OWL 2 Rule (RL) reasoning, allowing us to use off-the-shelf reasoning on top of our domain ontology. It is worth mentioning, however, that since our framework capitalises on existing, well-known standards (RDF, OWL, SPARQL, etc.), it is interoperable and it does not depend on specific implementations. For example, it requires minor updates to migrate to different triple stores, according to the application requirements, such as Strabon and AllegroGraph, or to use different SPARQL query engines.

F. Interlinking

To exploit the wealth of data, there comes the need of generating intelligent interconnections between different datasets. In literature, many systems have been implemented dealing with this issue but due to the vast heterogeneity of data, using existing systems into new datasets does not work in most cases. Interlinking is achieved based on geospatial data characteristics in some cases [4], while in others specific mechanisms have been developed to meet the needs of the data [41].

The system that is presented in [4] receives data from heterogeneous sources such as meteorological, health and earth observation. It creates an appropriate RDF representation and associations between specific characteristics. Data interlinking is achieved by calculating the similarity between different datasets. In [41] a system that integrates earth observation data to support data management is introduced. The system receives data from different data sources and has two different functionalities. For data that are related with China multiple components have been developed to adapt in different interfaces, while for international data the GEO DAB agent is used. PREDICAT [28] is a system that focuses on natural catastrophes prediction. PREDICAT uses different ontologies to semantically represent the data that are pertinent to the system (semantic layer). The system overcomes data heterogeneity and provides a common structure of interconnected objects containing spatiotemporal information (data integration layer), while a reasoner and a decision maker are implemented to provide the appropriate responses to the user (data processing layer). In CANDELA project²⁰ semantic

¹⁰<https://www.sqlshack.com/understanding-benefits-of-graph-databases-over-relational-databases-through-self-joins-in-sql-server/>

¹¹<http://graphdb.ontotext.com/>

¹²<https://rdf4j.org/>

¹³<https://virtuoso.openlinksw.com/>

¹⁴<http://ontop-spatial.di.uoa.gr/>

¹⁵<https://www.oracle.com/database/technologies/spatialandgraph.html>

¹⁶<https://allegrograph.com/>

¹⁷<https://www.stardog.com/>

¹⁸<https://www.openhub.net/p/useekm>

¹⁹<https://github.com/SemWebCentral/parliament>

²⁰<http://candela-h2020.eu/>

search is supported on EO images and other associated meta-data, using existing technologies, namely GeoSPARQL, OWL-Time, SOSA, DCAT and PROV-O. These technologies support a monitoring use case in agriculture, where the impact of a storm on vineyards is measured by first extracting knowledge from Sentinel images and then semantically fusing them with weather reports in the same area of interest²¹. However, the semantic search targets mainly to offer a mature solution for insurance companies, while it is not straightforward to monitor CAP-related regulations that require to be checked in terms of the farmers' compliance to CAP rules.

G. The complete life-cycle for CAP monitoring using semantic technologies and linked open EO data

Despite the fact that a lot of progress has been achieved in different aspects of the individual components mentioned above (e.g. Semantic annotation has been implemented under the environmental monitoring domain [27] [28] [42], data integration has been widely investigated (Interlinking) [33] [43] [44] [45] [28] [46]), not much effort has been achieved in implementing a system that supports the **complete life-cycle**. The major challenge, which we address in this paper, is mostly related with interlinking phase and dealing with the heterogeneity of data (e.g. sensors [46], aerial and satellite imagery [43], OpenStreetMap data [43] [45], Google Earth imagery [45]) and defining the ways to exploit these data and enhance knowledge discovery [33].

Contrary to the presented approaches in this Section, in this work we focus on the linked open EO data life cycle paradigm proposed in [3], aiming to support impactful use cases in CAP. To the best of our knowledge, this is the first attempt to reuse and adapt the proposed architecture to the domain of CAP monitoring. The proposed framework implements a hybrid scheme of data analysis and annotation: the results of a data-driven crop classification framework are semantically annotated and interlinked in order to foster advanced interpretation, such as improving classification accuracy through domain knowledge, and querying solutions. We demonstrate the added value and feasibility of our approach in a number of challenging use cases in CAP monitoring.

III. METHODOLOGY

The overall framework of our proposed methodology is presented in Figure 2. The layers are the image analysis layer, the mapping layer, the data ingestion and reasoning layer and the query processing layer. The knowledge extraction phase of the life-cycle, which was presented in Section II, consists of the image analysis layer using machine learning techniques for the content extraction. The semantic web technologies in the context of EO and Agriculture domains involve also the semantic annotation and transformation into RDF of the Mapping layer. The data ingestion and reasoning layer populates the knowledge base with the extracted knowledge, for storage and querying in a standard data representation model. Finally, the interlinking is done as part of the query

processing layer that allows for checking the compliance of the farmers' declarations to the CAP regulations under the reasoning mechanism of the previous layers, using the open linked data paradigm. These layers are presented in detail in the following sections.

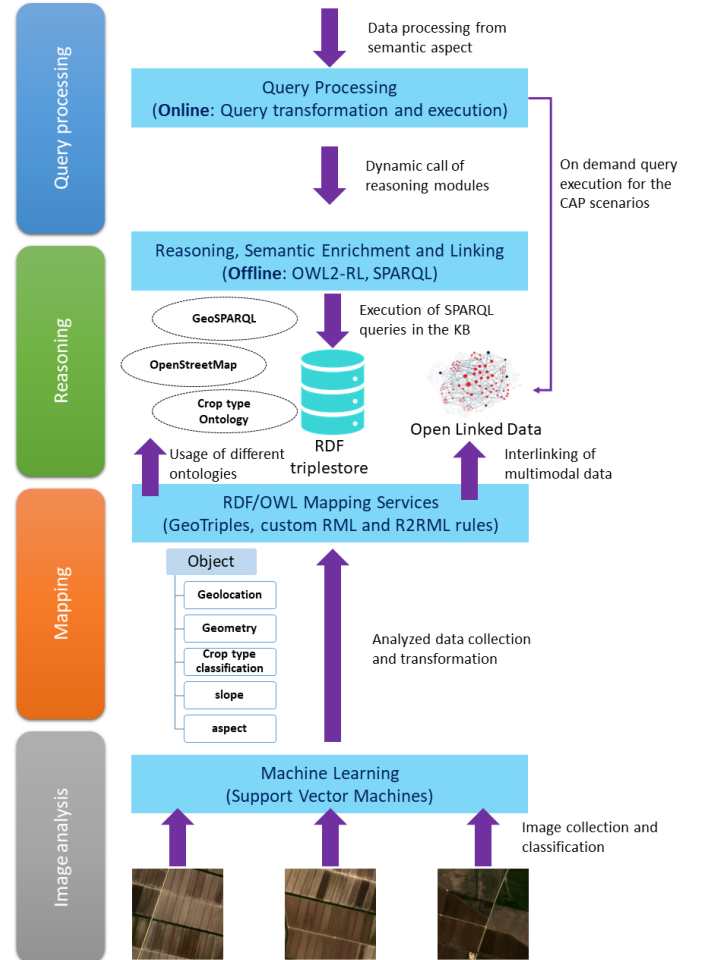


Fig. 2: System architecture overview

A. Satellite image analysis for the monitoring of the CAP

Paying agencies of EU Member States, usually receive the annual subsidy applications in May or June. The paying agency inspectors require the information of the cultivated crop type, even as early as May. This way, inspectors can select and organize their OTSCs, which follow in the coming months. Additionally, crop classification results, received prior to the annual farmer declarations, can assist as an alerting mechanism during the application process. The image analysis layer is the first layer of our proposed pipeline (Figure 2), where a Sentinel-based crop classification system for the monitoring of the CAP is developed.

The Area of Interest (AOI) is located in northeastern Spain and specifically the district of Navarra. The AOI covers the agricultural land surrounding the city of Pamplona, capital of Navarra. In detail, the dataset includes 9,052 parcels and amounts to approximately 215 km² of total land area. The Northern part

²¹<http://candela-h2020.eu/content/semantic-search-v2>

of Navarra is surrounded by the Pyrenees Mountains, as they stretch southward from France. The landscape of the district is a mixture of forested mountains and watered valleys, while the agricultural land is characterized by substantial fragmentation [47]. This study builds upon the crop classification results of [8], as described in section II. SVM based crop maps are produced, including the crop types of soft wheat (50%), barley (26%), oats (8.4%), maize (1.4%), sunflower (3.2%), vineyards (1.3%), broad beans (4.5%), rapeseed (5.4%) and cherry trees (0.2%). The aforementioned crop types are the lowest level of ontology, as shown in Figure 5.

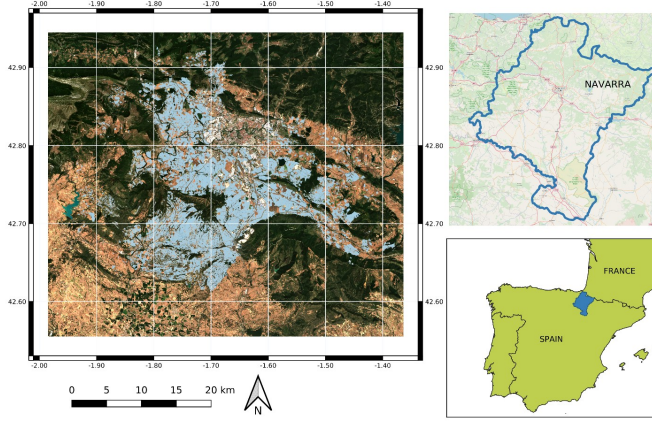


Fig. 3: Study area located in northeastern Spain and specifically in the Navarra district. The parcels of interest are shown in light blue color.

The dataset for training the SVM classifier is based on the Land Parcel Identification System (LPIS), which includes the parcel polygons in vector format and the associated farmer declaration for the 2018 CAP subsidy applications. The parcel polygons are used for segmenting the stack of Sentinel imagery to objects. The LPIS was provided by INTIA²², a public company, part of the Department of Rural Development, Environment and Local Administration of Spain. INTIA serves the role of paying agency for the district of Navarra, performing all CAP compliance inspections for the area. INTIA has additionally provided the timeline of growth for the major crops of the area. Figure 4 illustrates the acquisitions of Sentinel-2 images, spanning over the entirety of the various crop cycles. The feature space used for the crop classification includes the Sentinel-2 images for the acquisitions depicted in Figure 4. The acquisitions have been selected to have minimal cloud coverage over the AOI. All spectral bands, except B09 and B10, were used, along with the Vegetation Indices (VI) Normalized Difference Vegetation Index (NDVI), Normalized Difference Water Index (NDWI) and Plant Senescence Reflectance Index (PSRI). In this study, NDWI is used as defined by Gao (1996) [48]. Sentinel-2 images are atmospherically corrected to Bottom Of Atmosphere (BOA) reflectances using the Sen2Cor tool, and all bands are resampled to 10 m spatial resolution. The feature space comprises of parcel entities de-




Fig. 4: Timeline of the growth cycle of major crops in Navarra, reworked from [8], together with the acquisition dates of the Sentinel-2 images.

scribed by the mean value, for all features, of the pixels that fall within their LPIS boundaries.

The proposed methodology is based on a traffic light system approach. Specifically, each parcel is categorized into four groups, each offering different levels of confidence. These categories comprise of the green, yellow, red and unreliable classes, indicating high to low levels of confidence, in that order. The categorization of each parcel is based on the difference between the two highest SVM scores. This study focuses predominantly on the green category, namely the decision of highest confidence. It is on those most confident samples that we then record the mismatches of model predictions and the farmer declarations. An alarm mechanism is then introduced, identifying the green parcels that have been systematically misclassified (mismatch of declaration and prediction) during the cultivating season.

In the proposed algorithm (Algorithm 1), the alarms of potential false declarations are detected for any time instance throughout the year, with variable accuracy considering the satellite imagery available to date. The X_{train} and X_{test} are the training and test feature spaces, respectively. The feature spaces are dynamically populated with all new acquisitions. Therefore, when the algorithm is executed (*currentDate*), it uses the up to date feature spaces as input; containing imagery until the latest available acquisition (Equation 1).

$$acqDate = acqDate_k, k = 1, \dots, \lambda, \dots, A \quad (1)$$

where A is the index to the latest acquisition prior to *currentDate* and λ is the index to the acquisition that defines the starting feature space, early in the year.

The algorithm iterates $A - \lambda$ times, each time recording the misclassifications (*mis*). Misclassifications, in this context and as previously stated, refer to the mismatch between the SVM model's prediction and the farmers' declaration. For each iteration, a 2nd order polynomial Support Vector Machine (SVM) model is trained based on $X_{train}(t)$ (Equation 2) and the SVM scores are computed after applying the model to $X_{test}(t)$ (Equation 3). $X_{train}(t)$ and $X_{test}(t)$ are the training and test data for each iteration, and f is the number of individual features for each acquisition. The farmer declarations, as part

²²<https://www.intiasa.es/en/>

Algorithm 1 Smart Sampling

Input: $X_{train} = \{(x_i, d_i), x_i \in R^l, d_i \in \{1, 2, \dots, m\}, i = 1, \dots, N\}$, $X_{test} = \{(x_i, d_i), x_i \in R^l, d_i \in \{1, 2, \dots, m\}, i = 1, \dots, M\}$, $acqDate = \{acqDate_k, k = 1, \dots, \lambda, \dots, A\}$, number of iterations $t = 0$, the acquisition number that bounds that starting feature space λ , the date the algorithm is executed $currentDate$, persistence threshold for each iteration $P_t = 0$, number of misclassifications $mis = \{mis_i, i = 1, \dots, M\}$

Output: Alarms for potential breaches of compliance

```

while  $acqDate_t \leq currentDate$  do
   $X_{train}(t) = \{(x_i, d_i), x_i \in R^{(acqDate_{(t+\lambda)} \cdot f)}, d_i \in \{1, 2, \dots, m\}, i = 1, \dots, N\}$ 
   $X_{test}(t) = \{(x_i, d_i), x_i \in R^{(acqDate_{(t+\lambda)} \cdot f)}, d_i \in \{1, 2, \dots, m\}, i = 1, \dots, M\}$ 
   $h(\vec{x}) = \sum \alpha_i y_i (\vec{x}_i \cdot \vec{x} + b)^2 + b$  (Train SVM)
   $\pi_d = P(d|X_{test}(t)) = \frac{1}{(1 + e^{(A * f(X_{test}(t)) + B)})}$  (Calculate SVM Scores)
   $score = \max \pi_d - \max(\pi_d - \max \pi_d)$ 
  if  $t \bmod 2 = 1$  then
     $P_t = P_t + 1$ 
  end if
   $alarms = \{\}$ 
  for  $i = 1$  to  $M$  do
    if not  $score \geq threshold$  then
      continue (Bypass unreliable decisions)
    end if
    if parcel is misclassified then
       $mis_i = mis_i + 1$ 
    end if
    if  $mis_i \geq P_t$  and  $mis_i > 0$  then
       $alarms = alarms \cup i$ 
    end if
  end for
   $t = t + 1$ 
end while

```

of the annual subsidy application for the CAP, are used for labeling the parcels and thus training the model. A stratified random split was performed to split the samples into 30% and 70% subsets for $X_{train}(t)$ and $X_{test}(t)$ respectively. This amounts to 2,716 parcels, which have been used for training. All classification metrics that are presented in later sections have been averaged for 20 random splits of different seeds. The percentage of training samples was ultimately set to 30% after experimenting with larger datasets, which have provided only a marginal increase in performance.

$$X_{train}(t) = \{(x_i, d_i), x_i \in R^{(acqDate_{(t+\lambda)} \cdot f)}, d_i \in \{1, 2, \dots, m\}, i = 1, \dots, N\} \quad (2)$$

$$X_{test}(t) = \{(x_i, d_i), x_i \in R^{(acqDate_{(t+\lambda)} \cdot f)}, d_i \in \{1, 2, \dots, m\}, i = 1, \dots, M\} \quad (3)$$

where x_i is the feature representation of the i -th out of N parcels, belonging to $R^{(acqDate_{(t+\lambda)} \cdot f)}$. The superscript represents the dimensionality of the feature space. In each iteration, starting with $t = 0$, the feature space comprises of the starting feature space, i.e. the one including all features, f , of all acquisitions until $acqDate_\lambda$, plus all features, f , of acquisitions $acqDate_{\lambda+t}$; d_i is the label for each parcel, ranging from 1 to m ($=10$), representing the different crop types.

The difference between the two highest per class scores $P(d|X_{test}(t))$ for each sample is recorded as the overall score value (Equation 4) for the selection of the most confident decisions against a *threshold*. These parcels constitute the *green* labels in the aforementioned defined traffic light system. We denote by π_d the difference scores, i.e.: $\pi_d = P(d|X_{test}(t))$, so the overall score is given by:

$$score = \max \pi_d - \max(\pi_d - \max \pi_d) \quad (4)$$

The algorithm returns the misclassifications of the last iteration, namely the confident decisions of mismatch between the prediction and the declaration, which are classified as *alarms*. *alarms* are the samples that have been found misclassified at least P_t times, in all previous iterations. P_t is varying based on the time within the year the algorithm is executed (Equation 5).

$$P_t = \sum_{t=1}^{A-\lambda} t \bmod 2 \quad (5)$$

Early classifications are characterized by limited reliability, as the imagery included the training datasets does not cover the entirety of crops' growth cycle. For this reason Algorithm 2 can be optionally used to further refine the selected alarms.

Algorithm 2 Smart Sampling Fine Tuning

Input: *alarms* from Algorithm 1, $Y_{test} = \{(y_i \in \{1, 2, \dots, m\})$ the actual estimations of the classifier, $D_{test} = \{(d_i \in \{1, 2, \dots, m\})$ the declared labels

Output: Updated alarms for potential breaches of compliance

```

 $updatedAlarms = \{\}$ 
 $i = 1$ 
 $n = \text{size}(alarms)$ 
while  $i \leq n$  do
  if season of  $y_i \neq$  season of  $d_i$  then
     $updatedAlarms = updatedAlarms \cup i$ 
  end if
   $i = i + 1$ 
end while

```

Algorithm 2 uses the alarms of Algorithm 1 as input and returns an updated set of alarms. In order to increase the reliability of the smart sampling algorithm, we select alarms

for which the crops are classified to a type of a completely different crop season class that then one of the type declared.

B. Mapping layer for semantic representation

An important aspect of the framework is the representation of the available information, e.g. crop classification results, as well as capturing of domain knowledge needed to further correlated results. For the former, we use GeoTriples to transform data into the RDF format, while for the later we developed a domain-specific ontology. The metadata and analysis results are integrated using geotagged information and analyzed using semantic queries.

As far as the domain ontology is concerned, it consists of three layers:

- I Data, which describe crop land taxonomy. Figure 5 shows the relationship between period, crop family, crop type and crop code classes. More specifically, each crop type (crop type classes) has a unique crop code (crop code property) and belongs to a specific crop family (family classes). Crop families are connected with the season that each crop thrives (period classes). Each crop thrives in a different season, while some crops seem to thrive in year-round basis. Same crop families may belong to different seasons when they represent different crop types.
- II The crop type classification data that contain a linking of metadata which describe field information like parcel identifier, geometry, slope, aspect and classification scores for all different crop types.
- III Data collected from OpenStreetMap, containing water and waterways information in geospatial format i.e. the geometries of hydrographic network objects.

Figure 5 depicts the relationship between crop fields, periods, families, types and codes. Crop field is a class that corresponds to the crop fields that have been identified in a classification run. Each crop type is a subclass of a specific family and related to a specific crop code. All periods, families and types are different classes in the ontology. Crop codes are the main interconnection point between the crop fields and the characteristics of each crop type. OWL property restrictions have been identified to automatically detect the period of the crop code declaration or classification using the value of the crop code. These relationships are useful for extending the dataset described in layer II of the ontology (crop type classification data). For instance, in the example presented in Listing 2 the crop code that has been identified by classification is 1082 and the crop field is associated with the classification period winter. The crop code that the farmer declared is 1334, whose classification period is year-round. OWL2 RL is used to make useful inferences and apply reasoning rules into crop type classification data.

```

:SoftWeat
  a owl:Class ;
  rdfs:subClassOf :Cereals ;
  rdfs:subClassOf [
    a owl:Restriction ;
    owl:hasValue :winter ;
    owl:onProperty :hasClassificationPeriod ;
  ] ;
  owl:equivalentClass [
    a owl:Restriction ;

```

```

    owl:hasValue "1082" ;
    owl:onProperty vocab:has_CROPTYPE ;
  ] ;

```

Listing 1: Example OWL2 RL

In the example presented in Listing 1, the classification period and hyperclass have been automatically assigned to the crop field, taking advantage of the dynamics of OWL2 RL. T-Box reasoning (in the form of OWL 2 RL entailment rules supported by the GraphDB implementation) is applied to infer that soft wheat is a cereal, crop type “1082” corresponds to soft wheat and classification period “winter”, etc. More specifically, the `rdfs:subClassOf` has been used to assign the crop family values, while `owl:Restrictions` have been used to assign the value “winter” on the property `hasClassificationPeriod` when the value of the property `vocab:has_CROPTYPE` is “1082”.

```

@prefix : <http://mklab.it/ontologies/croptypes/> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix vocab: <http://example.com/ontology#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix fa: <http://example.com/farmer_ontology#> .
@prefix map: <http://example.com#> .
@prefix geo: <http://www.opengis.net/ont/geosparql#> .

<http://example.com/parcels_classification_wscores_v3/Geometry/14356>
  a geo:Geometry ;
  geo:asWKT
    "<http://www.opengis.net/def/crs/EPSG/0/4326>
    MULTIPOLYGON (((-1.6927107344688275 42.649935712372795,
    ... -1.692387508984902 42.65003808978244,
    -1.692421288894873 42.65001269177615,
    -1.6925231315390605 42.649985701208124,
    -1.6927107344688275 42.649935712372795)))
    )"^^geo:wktLiteral .

<http://example.com/parcels_classification_wscores_v3/id/14356>
  a
    fa:hasOwner vocab:parcels_classification_wscores_v3 ;
    vocab:has_ASPECT fa:farmland5 ;
    vocab:has_CROPTYPE 5.33798E1 ;
    vocab:has_CROP_CODE_ "1082" ;
    vocab:has_ID 1334 ;
    vocab:has_SLOPE 773280 ;
    vocab:has_SLOPE 5.0291E0 ;
    vocab:has_scores_t_1 7.511E-3 ;
    vocab:has_scores_t_2 4.33932E-1 ;
    vocab:has_scores_t_3 5.4892E-2 ;
    vocab:has_scores_t_4 1.41E-2 ;
    vocab:has_scores_t_5 4.9191E-2 ;
    vocab:has_scores_t_6 3.6199E-2 ;
    vocab:has_scores_t_7 2.17041E-1 ;
    vocab:has_scores_t_8 2.802E-2 ;
    vocab:has_scores_t_9 3.8657E-2 ;
    vocab:has_scores_typ 1.20457E-1 ;
    :hasClassificationPeriod :winter ;
    :hasDeclarationPeriod :year-round ;
    geo:hasGeometry
      <http://example.com/parcels_classification_wscores_v3/Geometry/14356> .

```

Listing 2: Example of crop field information as described in RDF combining the crop type ontology (I) and the results of crop type classification (II)

The transformation of crop type classification data and data collected from OpenStreetMap into the RDF model has been supported by GeoTriples. The tool accepts data in shapefile format and automatically produces an RDF Mapping Language (RML) file containing the rules that the RDF file should satisfy. Then, it produces the RDF serialisation that satisfies the RML rules. The Well-Known Text (WKT) representation of coordinate reference systems standard has been reused to capture location-related information.

For the classification results, GeoTriples tool is used multiple times to convert the results of each run into RDF format. Each file has a unique name and because of that, new instances are created in the Knowledge Base. In the end of this procedure, the Knowledge Base contains many different instances of the same parcel having in common the parcel identifier. In such way we keep crop type classification data

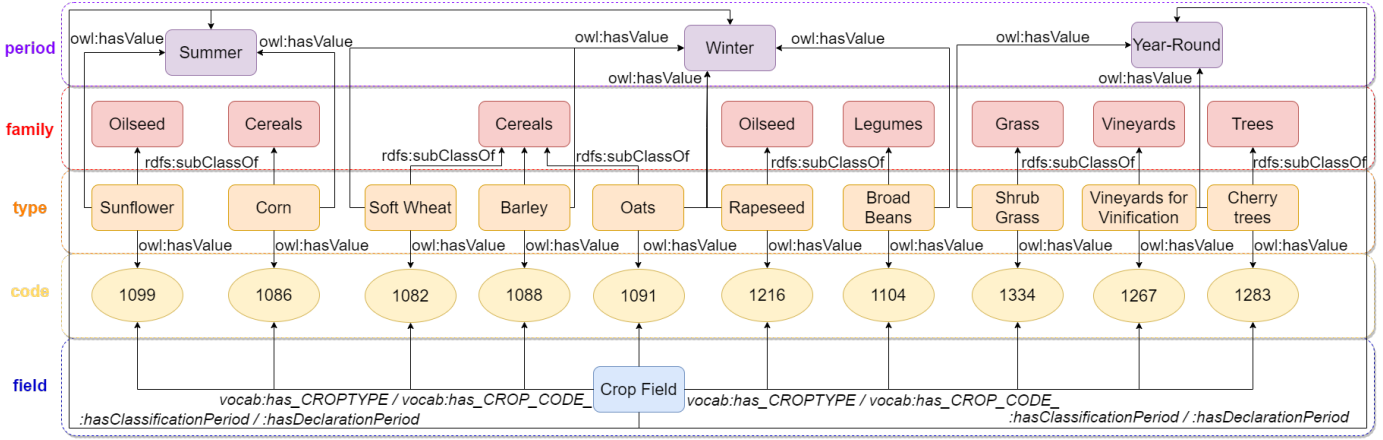


Fig. 5: Crop type ontology (knowledge extracted from [8]).

of past runs in the Knowledge Base, which can be aligned with new classification data.

C. Data ingestion and querying layer for storage and semantic enrichment

Semantic enrichment aims at interconnecting and further enriching the contents of the generated knowledge graphs, applying semantic rules. The focus is given on improving the smart sampling methodology, executing a set of queries (rules) to improve the selection process of OTSCs. In the following, we describe the specifics of the approach, presenting the defined rules.

With every new image acquisition a new crop classification is performed and the classification results are dynamically populating the knowledge base. Procedural code is used to run semantic queries in sequence and pass the needed values from past queries (parcel, value1, value2) into next ones (Listings 3 to 8). The reason for this decision is that SPARQL lacks in terms of arguments saving or passing into next queries and these calculations are better expressed using many SPARQL queries, improving the execution time. The queries that are presented in this section are running for each parcel instance. The parcel instances are retrieved by the query depicted in Listing 3. A parcel instance (described as *<parcel>*) may be for example: http://example.com/parcels_classification_wscores_v8/id/1109. The first step (Listing 4) is to retrieve the two highest classification scores among the 10 different crop types, as these are defined in the ontology (Figure 5). This is done in order to compute the *score* value in Algorithm 1. This process takes place for each of the different classification instances in the knowledge base. In this query, typical SPARQL functions are used, such as BIND, to group different type results under the same variable, and ORDER BY to arrange the results in descending order.

```
PREFIX vocab: <http://example.com/ontology#>
SELECT * WHERE {
  ?parcel a vocab:parcels_classification_wscores_v8.
}
```

Listing 3: Semantic query to retrieve all parcel instances for the latest classification run

```
PREFIX vocab: <http://example.com/ontology#>
select ?max_types where {
  <parcel> vocab:has_scores_t_1 ?type_1.
  <parcel> vocab:has_scores_t_2 ?type_2.
  <parcel> vocab:has_scores_t_3 ?type_3.
  <parcel> vocab:has_scores_t_4 ?type_4.
  <parcel> vocab:has_scores_t_5 ?type_5.
  <parcel> vocab:has_scores_t_6 ?type_6.
  <parcel> vocab:has_scores_t_7 ?type_7.
  <parcel> vocab:has_scores_t_8 ?type_8.
  <parcel> vocab:has_scores_t_9 ?type_9.
  <parcel> vocab:has_scores_t_10 ?type_10.
  {BIND(?type_1 as ?max_types)}
  union
  {BIND(?type_2 as ?max_types)}
  union
  {BIND(?type_3 as ?max_types)}
  union
  {BIND(?type_4 as ?max_types)}
  union
  {BIND(?type_5 as ?max_types)}
  union
  {BIND(?type_6 as ?max_types)}
  union
  {BIND(?type_7 as ?max_types)}
  union
  {BIND(?type_8 as ?max_types)}
  union
  {BIND(?type_9 as ?max_types)}
  union
  {BIND(?type_10 as ?max_types)}
}
ORDER BY DESC (?max_types)
LIMIT 2
```

Listing 4: Semantic query to retrieve the two highest classification score values per parcel

The difference between the two largest values (Listing 5) is calculated using the results of the previous query (value1, value2). If the difference is bigger than a specific threshold, the query in Listing 6 is used to mark the parcel as “green”, i.e. parcels with a high confidence that the prediction is correct.

```
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT * where {
  BIND (xsd:double(value1)-xsd:double(value2)
  AS ?result)
  FILTER (?result > "0.5"^^xsd:double)
}
```

Listing 5: Semantic query to detect the parcels that the difference from the two highest score values is above threshold

```
PREFIX vocab: <http://example.com/ontology#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX map: <http://example.com/#>
```

```

INSERT {
  <parcel> map:risk "green parcel".
} WHERE {
  <parcel> ?p ?o.
}

```

Listing 6: Semantic query to mark green parcels that satisfy the conditions of the two previous queries

In addition, each time a green parcel is identified, the query in Listing 7 marks the parcels that have been incorrectly classified. In the end, the query in Listing 8 marks the parcels that the season of the declaration does not agree with the season of the classification, based on the domain ontology (Figure 5). SPARQL INSERT function is used to enrich the parcels with “misclassification”, “green parcel” and “same season” information. All in all, semantic enrichment tries to enrich smart sampling, taking into account the provided classification results and domain knowledge about crop types, so as to detect parcels that have a high probability to have a false crop type declaration by the farmers.

```

PREFIX vocab: <http://example.com/ontology#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX map: <http://example.com/#>
INSERT {
  <parcel> map:misclassification "misclassification".
} WHERE {
  <parcel> map:risk "green parcel".
  <parcel> vocab:has_CROP_CODE_ ?decl .
  <parcel> vocab:has_CROPTYPE ?class .
  FILTER(?decl!=xsd:integer(?class))
}

```

Listing 7: Semantic query to mark parcels as misclassified when parcels have been marked as green and declaration does not agree with classification

```

PREFIX vocab: <http://example.com/ontology#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX map: <http://example.com/#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX : <http://mklab.iti.gr/ontologies/croptypes/>
INSERT {
  <parcel> map:same_season "false".
} WHERE {
  <parcel> vocab:has_ID ?id.
  <parcel> vocab:has_CROP_CODE_ ?decl .
  <parcel> vocab:has_CROPTYPE ?class .

  ?type a owl:Class.
  ?type rdfs:subClassOf ?object1 .
  ?object1 owl:hasValue ?code.
  ?type rdfs:subClassOf ?object2 .
  ?object2 owl:hasValue ?period1.
  ?period1 a :Period.
  FILTER regex (str(?code), str(?decl))

  ?type2 a owl:Class.
  ?type2 rdfs:subClassOf ?object3 .
  ?object3 owl:hasValue ?code2.
  ?type2 rdfs:subClassOf ?object4 .
  ?object4 owl:hasValue ?period2.
  ?period2 a :Period.
  FILTER regex (str(?code2), ?class)
  FILTER (?period1!=?period2)
}

```

Listing 8: Marking parcels where declaration and classification belong to different seasons

D. Query processing layer for interlinking spatial data

By capturing data in the RDF model space, we enable spatial relationships-based querying and easy integration with

other data sources, such as Linked Data. We demonstrate the query answering capabilities of the framework, as well as the ability to integrate external datasets, by defining queries to detect possible non-compliance of the farmers according to the specified rules.

1) *Greening 1 requirement:* The query in Listing 9 calculates the number of different crop types that a farmer cultivates and the total area of their farm. More specifically, the query detects a breach of compliance when the farmers own a total farm area between 10 and 30 hectares and cultivate at least two different crop types, but the dominant one is more than 75% of the total farm area. The crop type as it is calculated in the SVM classification is expressed in the mapping via the vocab:has_CROPTYPE property. In the RDF space, we use count(distinct ?ctype) as ?count to detect the number of the different crop type values that are detected for each farmer (GROUP BY ?owner). Additional queries are applied to also check, for instance, the farmers that have less than three different crop types cultivated in a total farm area of more than 30 hectares if: $\{sum > 300000\} \wedge \{count < 3\} \wedge \{max_{croptype}[1] + max_{croptype}[2] > 0.95 * sum\}$ as it has been described in Section I. The query (Listing 10) further supports the greening 1 requirement by detecting the farmers that grow less than 3 different crop types in a total area of more than 30 hectares and farmers that grow less than 2 different crop types in a total area of 10 to 30 hectares.

```

PREFIX geo: <http://www.opengis.net/ont/geosparql#>
PREFIX ext: <http://rdf.useekm.com/ext#>
PREFIX fa: <http://example.com/farmer_ontology#>
PREFIX vocab: <http://example.com/ontology#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
select * where {
  {
    select ?owner (sum(?area) as ?max) ?sum where {
      { select ?owner (count(distinct ?ctype) as ?count)
        (sum(?area) as ?sum) where {
          ?field fa:hasOwner ?owner .
          ?field vocab:has_CROPTYPE ?ctype .
          ?field geo:hasGeometry ?geo.
          ?geo geo:asWKT ?polygon.
          BIND(ext:area(?polygon) as ?area).
          ?owner a fa:Farmer .
        }
      }
      GROUP BY ?owner
      HAVING (?sum > 100000 && ?sum <= 300000 && ?count >=2)
    }
    ?field fa:hasOwner ?owner .
    ?field vocab:has_CROPTYPE ?ctype .
    ?field geo:hasGeometry ?geo.
    ?geo geo:asWKT ?polygon.
    BIND(ext:area(?polygon) as ?area).
    ?owner a fa:Farmer .
  }
  GROUP BY ?owner ?ctype ?sum
  ORDER BY DESC (?max)
}
FILTER (?max>0.75*?sum)

```

Listing 9: Semantic query to extract possible non-compliance in the Greening 1 requirement for the farmers owning a total farm area between 10 and 30 hectares and cultivating at least two different crop types, where the dominant one is more than 75% of the total farm area

```

PREFIX geo: <http://www.opengis.net/ont/geosparql#>
PREFIX ext: <http://rdf.useekm.com/ext#>
PREFIX fa: <http://example.com/farmer_ontology#>
PREFIX vocab: <http://example.com/ontology#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

```

```

select ?owner (count(distinct ?ctype) as ?count)
  (sum(?area) as ?sum) where {
  ?field fa:hasOwner ?owner .
  ?field vocab:has_CROPTYPE ?ctype .
  ?field geo:hasGeometry ?geo .
  ?geo geo:asWKT ?polygon .
  BIND(ext:area(?polygon) as ?area) .
  ?owner a fa:Farmer .
}
GROUP BY ?owner
HAVING ((?sum > 300000 && ?count < 3) ||
  (?sum > 100000 && ?sum <= 300000 && ?count < 2))

```

Listing 10: Semantic query to extract possible non-compliance in the Greening 1 requirement for the farmers that grow less than 3 different crop types in a total area of more than 30 hectares and farmers that grow less than 2 different crop types in a total area of 10 to 30 hectares

The area of the fields is calculated using GeoSPARQL ext:area function, taking advantage of the polygon points coordinates. The number of different crop types is calculated using the result of the SVM classification prediction of crop type on the field, while the sum of the fields area using the results of GeoSPARQL area calculations. All operations are implemented using SPARQL functions such as count, distinct and sum.

2) *SMR 1 requirement:* Another important information for the end users is the distance of the parcels from the hydrographic network objects. This is in accordance with the requirements of SMR1, as described in Section I. In order to effectively identify the parcels susceptible to contribute nitrate-rich soil to nearby surface water, a filtering mechanism takes place, accounting for the slope and aspect of the parcel. Results are important for both farmers and paying agencies that want to check the compliance according to SMR 1 requirement.

Listing 11 presents the query that lists the parcel instances where the distance from surface waters is lower than 10 meters, which is the buffer for organic manure application. Since both parcel (LPIS) and hydrographic network data²³ (data from OpenStreetMap) contain geospatial information (e.g. multipolygon, polygon etc.), the distance is calculated using geof:distance function of GeoSPARQL. The function accepts two geometries and calculates the shortest distance between any two points of the specified geometries. A filtering mechanism selects the fields according to their distance from the hydrographic network objects, slope and aspect as described in section IV-C3. The GeoSPARQL function ext:closestPoint is used to compute the closest points of each geometry compared to the other geometry. A string replacement pattern is used in order to retrieve the coordinates of the two points, which are utilized to compute the angle of the two points in degrees. GraphDB math functions are used to achieve such computations.

```

PREFIX geof: <http://www.opengis.net/def/function/geosparql/>
PREFIX geo: <http://www.opengis.net/ont/geosparql#>
PREFIX uom: <http://www.opengis.net/def/uom/OGC/1.0/>
PREFIX map: <http://example.com/#>
PREFIX ogc: <http://www.opengis.net/ont/geosparql#>
PREFIX ofn: <http://www.ontotext.com/sparql/functions/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX ext: <http://rdf.usackm.com/ext#>
PREFIX vocab: <http://example.com/ontology#>
select ?parcel ?distance where {
  BIND(<parcel> AS ?parcel) .
  ?parcel a vocab:parcels_classification_wscores_v8 .
  ?a vocab:has_ID ?id .

```

```

?parcel vocab:has_SLOPE ?slope .
FILTER(?slope > 12)
?parcel ogc:hasGeometry ?s .

?s geo:asWKT ?o .
?fGeom geo:asWKT ?fWKT .
?fGeom map:containsWater ?b .
FILTER (?fGeom != ?s) .
FILTER NOT EXISTS {
  ?s map:containsWater ?wa .
}
BIND(geof:distance(?o, ?fWKT) as ?distance) .

FILTER(?distance <= 0.1)
BIND(ext:closestPoint(?o, ?fWKT) as ?c1point1)
BIND(ext:closestPoint(?fWKT, ?o) as ?c1point2)

BIND( replace( str(?c1point1), "[0-9\\.|-]*([0-9\\.|-])+",
  ".$1" ) as ?long )
BIND( replace( str(?c1point1), "([0-9\\.|-])+", ".$1" ) AS ?lat )

BIND( replace( str(?c1point2), "[0-9\\.|-]*([0-9\\.|-])+",
  ".$1" ) AS ?long2 )
BIND( replace( str(?c1point2), "([0-9\\.|-])+", ".$1" ) AS ?lat2 )
BIND (xsd:double(?lat)-xsd:double(?lat2) AS ?x)
BIND (xsd:double(?long)-xsd:double(?long2) AS ?y)

BIND (ofn:atan2(?x, ?y) AS ?atan2)
?parcel vocab:has_ASPECT ?aspect .
BIND( IF (ofn:toDegrees(?atan2) > "0"^^xsd:double ,
  ofn:toDegrees(?atan2), ofn:toDegrees(?atan2)+
  "360"^^xsd:double ) AS ?toDegrees )
BIND( IF (?aspect - 45 > 0, ?aspect - 45, ?aspect + 360 - 45) AS ?min )
BIND( IF (?aspect + 45 > 360, ?aspect - 360 + 45, ?aspect + 45) AS ?max )

FILTER(?min < ?toDegrees && ?toDegrees < ?max)
}
LIMIT 1

```

Listing 11: Semantic query to detect susceptible parcels according to the SMR 1 requirement taking into account the slope, the aspect, the angle and the distance of the parcel from hydrographic network objects

3) *Smart sampling:* The query in Listing 13 supports the retrieval of the results of section III-C. When querying for the parcels to be inspected through OTSC, at any given time in the year, all past classification instances until that point are used. More specifically, the query takes advantage of all past classification decisions for each parcel to ensure that the prediction is indeed a misclassification. The threshold, above which a parcel is considered to be persistently misclassified, is defined in Listing 12 and is dynamically updated given the different parcel classes, that correspond to different times of the year, which exist in the triple store. The different classification runs are saved as instances of vocab:parcels_classification_wscores_v1 for the first run, vocab:parcels_classification_wscores_v2 for the second, etc. In such way, using the FILTER function, we select the number of different runs that currently exist in the triple store by detecting the number of different classes that contain the word “parcels”. The query in Listing 13 retrieves all the parcels that have been misclassified in past classifications more times than the defined threshold. Then the associated season for the crop type prediction of the latest available classification is compared with the associated season type of the declarations using the results of Listing 8. If there is a disagreement, the parcel is selected as an alert and candidate for OTSC. This query further narrows the smart sampling filter and is recommended for query executions early in the year (months before July), when the classification results are less trustworthy. More information about this functionality is described in Algorithm 2.

```

PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX ofn: <http://www.ontotext.com/sparql/functions/>

```

```

SELECT ?final_threshold WHERE {
  {
    SELECT (count(distinct ?parcel_class) as ?count)

```

²³<https://download.geofabrik.de/europe/spain.html>

```

WHERE {
  ?parcel_instance a ?parcel_class .
  FILTER contains(str(?parcel_class), "parcels")
}
}
BIND(IF(ofn:floorMod(?count,2)="1"^^xsd:int,
  ((?count+"1"^^xsd:int)/2)-"1"^^xsd:int,
  ?count/2) AS ?threshold )
BIND(IF(?count="1"^^xsd:double, "0"^^xsd:int,
  ?threshold) AS ?threshold )
BIND(IF(?threshold > "5"^^xsd:double, "5"^^xsd:int,
  ?threshold) AS ?final_threshold )
}

```

Listing 12: Semantic query to calculate the smart sampling threshold according to the number of different parcel classes, that correspond to different times of the year, which exist in the triplestore

```

PREFIX map: <http://example.com/#>
PREFIX vocab: <http://example.com/ontology#>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
SELECT ?id (count(distinct ?parcel) as ?count) WHERE {

  ?parcel vocab:has_ID ?id.
  ?parcel map:misclassification ?o .

  ?latest_parcel a vocab:parcels_classification_wscores_v8.
  ?latest_parcel vocab:has_ID ?id.
  ?latest_parcel map:same_season "false".
}
GROUP BY ?id
HAVING (?count > final_threshold)
ORDER BY desc (?count)

```

Listing 13: Semantic query to check persistent misclassifications using the smart sampling threshold and the number of past classification instances that have been misclassified provided that in the latest classification run the period of the farmer declaration does not agree with the period of classification

IV. EXPERIMENTS AND RESULTS

A. Hyperparameter optimization and performance of the SVM classifier

The crop classification model was built using the SVC function from the scikit-learn library²⁴ of Python. The hyperparameters were optimized using grid search over a range of values for each parameter, using 5-fold cross-validation. The hyperparameter combination that was selected was the one that produced the highest overall accuracy. For the penalty parameter of the error term, C , we tested values within the range $[2^{-2}, \dots, 2^9]$. For the kernel coefficient, namely the gamma parameter, the optimal value was examined within the range $[10^{-4}, \dots, 10^1]$; while for the independent term of the kernel function, coef0 , we searched the range $[10^{-3}, \dots, 10^2]$. The best combination consisted of $C = 4$, $\gamma = 0.001$ and $\text{coef0} = 10$.

Finally, the SVC function includes a parameter called *class_weight*, which was set to “balanced”. Its impact is that it sets the parameter C of each class i to $\text{class_weight}[i] * C$, where $\text{class_weight}[i]$ is inversely proportional to class frequencies of the input data. This way the negative effects of an unbalanced dataset, such as the one used for this study, are ameliorated.

We evaluated our core classification model [8] for three different agricultural areas, in Greece, Lithuania and Spain, for which we had independent in-situ validation data. These areas present various challenges. In the case of Greece, the agricultural landscape is significantly fragmented, resulting in small and narrow parcels, occupied by mixed pixels. In Lithuania, there is extended cloud coverage throughout most of the season, resulting in sparse time-series of cloud free Sentinel-2 imagery. Ten, eleven and fourteen crop classes have been classified in the Spanish, Greek and Lithuanian cases, respectively.

Validated results were consolidated based on OTSCs that were performed by the respective paying agencies of the three AOIs, during the 2018 subsidy applications. In the case of the Spanish AOI, out of the 107 randomly selected parcels for inspection, 105 were classified correctly. In Greece, inspectors visited only parcels classified with high confidence, namely of high posterior probability for the classification decision, to crop types other than the one declared. These instances are considered as potential breaches of compliance. It was shown that 76 out of 85 inspected parcels were indeed wrongly declared and correctly classified by our model. Finally, in Lithuania, the validated dataset acquired through the inspections resulted in an overall accuracy of 76.2% in late June and 80% in late August out of 3,319 parcels inspected. The results revealed the dependencies of the crop classification model performance, on the percentage of truthful declarations, the cloud coverage and the parcel shape and size.

In the Spanish AOI for which we focus in this work all these dependencies were optimal, i.e. more than 97% of truthful declarations, limited cloud coverage, and an average parcel size of 2 ha. Hence more than 90% classification accuracy was achieved.

Our model evaluation analysis also revealed that classification decisions for larger parcels and parcels with straighter borders tend to have higher accuracy than smaller parcels or parcels with more irregularly-shaped boundaries. The parcel area is important since accuracy depends on the number of image pixels that fall within the parcel boundaries. Sentinel-2’s 10 m pixel size equates to 50 image pixels in 0.5 ha of land. Our analysis shows that having 50 pixels of information provides accurate results, whereas for smaller parcels the decision is both less confident, namely of lower SVM score, and less accurate (Tables I and II).

TABLE I Accuracy of crop type classification for different parcel area ranges

| Spain | | | |
|--------|------------------|--------------------|-------|
| | | Accuracy % parcels | |
| Large | >1ha | 96.70 | 44.85 |
| Medium | 0.35ha < x < 1ha | 94.44 | 23.85 |
| Small | <0.35ha | 87.02 | 31.30 |

In Table II is observed that there is a strong correlation between the SVM score and accuracy. Indeed, the subset of parcels with an SVM score larger than 0.85, achieves an overall accuracy of more than 97%. On the other hand, the subset of parcels with SVM scores less than 0.7, achieves an overall accuracy of 66%.

²⁴<https://scikit-learn.org/>

TABLE II Relationship between SVM scores and overall accuracy

| Spain | | |
|----------------|----------|-----------|
| SVM score | Accuracy | % parcels |
| >0.85 | 97.88 | 77.43 |
| 0.7 < x < 0.85 | 89.46 | 10.27 |
| <0.7 | 66.32 | 11.30 |

B. Scenarios for the control of the CAP

1) *Scenario 1: Smart Sampling of OTSCs*: The paying agency inspectors search for parcels of farmers that have potentially falsely declared the cultivated crop type. The parcels prone to noncompliance are dynamically provided, starting from late June until the end of the cultivation period, to allow the inspectors to better target their inspections.

Farmer profile: Dan is a farmer. He has into his possession a field containing barley, while he has declared that the field's crop type is maize.

Analysis results: His field has been selected for OTSC because the crop classification has classified the cultivated crop type as barley, with high confidence. Even though the classification is not particularly trustworthy, being performed in early June, the parcel is marked as high risk by the smart sampling algorithm, as the prediction belongs to a different season class, namely winter.

Decision: The paying agency inspector needs to check the field, as there is a strong possibility that the farmer has wrongly declared the cultivated crop type.

2) *Scenario 2: Greening 1: Crop Diversification*: As previously stated, paying agencies need an automated system to detect non-complying farmers with respect to the Greening 1 requirement.

Farmer profile: Bob is a farmer. He has three fields in his possession covering an area of 27.4066 hectares. The total area of the fields that contain soft wheat is 27.1058 hectares, while for the crops that contain rapeseed the total area is 0.3008 hectares.

Analysis results: According to our analysis there seem to be two fields with soft wheat and one with rapeseed. The farmer does not comply with the rule because even though he cultivates two different crops, the total area of the soft wheat parcels exceeds the 75% of the total area of the fields that the farmer has into his possession.

Decision: The paying agency inspectors need to check these fields because there is a strong possibility that the farmer is not complying with the Greening 1 requirement.

3) *Scenario 3: detecting parcels prone to non-compliance to the SMR-1 requirement*: The paying agency inspectors need to establish the compliance to the SMR 1 defined buffer zones. In order to do that the distance from the parcels to nearby surface waters is calculated. The slope and aspect of the parcels are also taken into account in order to establish if there is an actual risk for runoff.

Farmer profile: Lucy is a farmer. She has into her possession a field whose distance from surface waters is almost 3 meters. The slope of the field is 15°, therefore of high runoff risk, and the aspect 251°, namely of western orientation. The aspect of the proximity line is 267°.

Analysis results: Since the parcel appears to be very close to surface waters, within the SMR-1 buffer, and the difference of the parcel and proximity line aspects is within the range of potential runoff, the analysis marks it as high risk.

Decision: The paying agency inspectors need to monitor this parcel because there is a strong risk for nitrate-rich runoff to surface waters.

C. Implementation of the CAP scenarios

1) *Implementation of Scenario 1: Smart Sampling*: In this scenario, the potential of better targeted, smart OTSCs is explored, exploiting accurate crop classification results, early in the year. This is achieved through the proposed semantic enrichment of classification results and the pertinent smart sampling query. The idea is to provide alarms, starting from early summer, when declarations are usually received, and dynamically update those using progressively larger feature spaces that include new Sentinel-2 acquisitions.

Table III shows the Producer's Accuracy (PA) and User's Accuracy (UA) for the predictions of all crop classes. PA is the percentage of correctly classified parcels against the total number of parcels of a given class. On the other hand, UA is the percentage of correctly classified parcels for a given class against the total number of parcels classified to that class. The metrics in Table III refer to the classification results produced using the entirety of acquisitions, 24/01/2018 to 21/10/2018, and have been averaged over 20 iterations of random training dataset splits. Additionally, the PA and UA results address the lowest level of crop taxonomy, as shown in Figure 5.

TABLE III PA and UA for the classification of 10 crops types using the full season time-series of Sentinel imagery

| Class | UA (%) | PA (%) |
|--------------|--------|--------|
| Soft wheat | 94.00 | 95.42 |
| Maize | 94.98 | 94.36 |
| Barley | 93.09 | 93.69 |
| Oats | 92.52 | 88.40 |
| Sunflower | 94.87 | 92.85 |
| Rapeseed | 95.34 | 92.08 |
| Broad beans | 94.84 | 89.40 |
| Shrub grass | 84.15 | 80.70 |
| Vineyards | 85.50 | 87.76 |
| Cherry trees | 83.22 | 81.98 |

Table III shows excellent classification performance for most of the crop types. Specifically, it is observed that the classifier achieves PA and UA values of more than 80% for all crop classes; while reaching values as high as 95%. It is also worth noting the excellent performance for the classes oats, soft wheat and barley. These crop types belong to both the cereals and winter superclasses, having similar spectral and phenological characteristics. Nonetheless, the classifier appears to discriminate among them easily. Finally, the weakest performances are observed for the classes of shrub grass, vineyards and cherry trees. These are all year-long crop types, with no significant phenological characteristics to assist the classifier. Additionally, the shrub grass class, which is an ambiguous crop description, is characterized by diverse spectral characteristics from parcel to parcel.

The metrics shown in Table III theoretically allow for the

effective sampling of OTSCs. However, the results presented have been produced using images until the end of October, when the notion of smart sampling becomes obsolete, as the inspections would have preceded. For this reason, the crop classification is performed at multiple instances throughout the year, starting as early as May 4. As it would be expected, the classification results of reduced feature spaces, when executed early in the year, would achieve suboptimal results of low reliability. In this regard, per class scores are computed for each sample, using class membership probability estimates [49]. Therefore each sample is associated with ten different membership probabilities, as many as the crop classes involved. In order to select the most reliable decisions, the difference between the two highest scores for each sample is recorded (*score* in Algorithm 1). This is achieved via the semantic enrichment mechanism, as described in figures 4 and 5. The following figure illustrates those differences in the form of a histogram, using 100 bins, for the entire dataset. The algorithm classifies most of the data with high reliability. The ones that are classified with low score are most likely crops that belong to the same family and some of them with similar spectral signatures are expected to be misclassified at some level.




Fig. 6: Histogram of differences between the two highest scores of each sample decision. The closer is the difference to 1.0 the lower the classification uncertainty.

The SVM algorithm is trained and tested for progressively larger feature spaces, whose classification results populate the knowledge base. The evolution of accuracy is shown in Figure 7, using the F1-score metric [50]. F1-score, defined as $2 \times (\frac{PA \times UA}{PA + UA})$, providing an overall accuracy metric that accounts for both PA and UA. The first point on the x-axis of Figure 7 refers to the F1-score for the classification that uses images until the 4th of May. For each next run onward, features keep on populating the feature space, with every new acquisition. It is observed that the larger feature spaces result in better classification results. Specifically, feature spaces that include the July 23rd acquisition, appear to approach optimal F1-scores for most crop classes, creating a plateau on the evolution curve. The majority of these crops are harvested from late June to early July, and thus our model can separate

the classes more comfortably, when features that cover this period are added to the feature space.




Fig. 7: Evolution of F1-score for crop type classification with progressively larger feature spaces

Accurate crop classification enables the monitoring of the CAP rules and allows for efficient decision making on the farmers' compliance. Towards this direction, the inspected parcels were assorted based on the previously described traffic light system. Two highest prediction probabilities, or SVM scores (Figure 8) are taken into consideration to pinpoint the parcels of the highest probability of noncompliance, for which the algorithm is assumed to predict with high reliability.

At first, for each run with a different feature space in the knowledge base, the set of green misclassifications is recorded, namely the reliable instances for which the prediction does not match the declaration. Then, the set of persistent green misclassifications is computed using the semantic query in Figure 13. Persistence refers to the number of times a sample has been found misclassified in the different classification iterations. In this study, there are two different thresholds of persistence. First, if a sample has been misclassified more than five times, in all different runs, from May to late October, then it is considered to be a validated alert. This is called constant persistent misclassification, as shown in Figure 9, and is assumed to function as the validation dataset, against which performance metrics are computed.

Figure 8 presents the number of constant persistent misclassifications relative to the increasing number of features, for different threshold values of what is considered a green parcel. The latter translates to varying values of the difference between the two highest scores for each sample. As expected, lower thresholds lead to higher number of misclassifications. It can also be observed that for each threshold that was tested, the pattern of the plot line is the same. The number of misclassifications is constant during May, then it presents an increase at the start of June, which begins to stabilize towards the start of August. It can be concluded that it is difficult for the algorithm to identify truly mislabeled data, really early in the year, but it seems to perform better with an increasing number of features. This improvement, which starts around June, is justified because most of the crops that were examined, are harvested at June, after which our model can classify the

data with higher confidence.




Fig. 8: Evolution of the number of constant persistent misclassifications for a varying threshold of the green category

However, in order to calculate the persistence of a parcel as mentioned above, a full Sentinel-2 series, for any given year of inspection, is required. Therefore, the constant persistent misclassifications are merely used as validation dataset of truly wrongly declared parcels, and are not part of the smart sampling algorithm. Thus, in order to be able to identify wrongly declared samples, in real scenarios, the total number of times that a green parcel has been misclassified, until the current run, is calculated. Then, we set a different threshold for each iteration (P_t in Algorithm 1) which is calculated using the semantic query in Figure 12. This is referred to as threshold of varying persistent misclassifications.

Figure 9 displays the number of misclassified parcels using both constant and varying persistence thresholds and a constant green threshold equal to 0.5. The threshold value was defined based on the a priori knowledge of the percentage of false declarations, annually. INTIA has stated that usually no more than 3% of the agricultural parcels in Navarra are falsely declared. Additionally, INTIA, as acting paying agency for the Navarra region, is obligated to conduct randomly selected inspections for at least 1% of applications. Therefore, using 0.5 as the threshold for the green category, provides 267 constant persistent misclassifications, which amount to 3.1% of the dataset; satisfying both the expected false declarations percentage and the minimum number of mandated OTSCs. Inspecting Figure 9, each point, of both curves, is associated with the persistence number that is used to count the misclassifications.

2) Implementation of Scenario 2: Crop Diversification:

Table IV contains a list of farmers that seem to not comply with the rule described in Figure 9 and Figure 10. The first column contains the farmer identifier that corresponds to a specific farmer, the second contains the count of different crop types that the specific farmer cultivates in the fields that he has into his possession and the third the sum area of the parcels that the farmer has into his possession. The fourth column contains the area of the main crop for the cases that the farmer cultivates more than 2 crop types in an area of 10 to 30 hectares but the main crop exceeds the 75% of the total area. It also contains the sum of the two main crops for the




Fig. 9: The evolution of the number of misclassifications under constant (blue line) and varying thresholds (orange line). The varying threshold increases by one every two image acquisitions. The thresholds define which parcels are the misclassifications for each classification instance looking at the number of times they were found misclassified in previous runs.

cases that the farmer cultivates more than 3 crop types in an area of more than 30 hectares but the sum of two main crops exceeds the 95% of total area. The last column contains the crop types that the farmer is cultivating in their farm.

TABLE IV List of farmers that seem to be non-compliant to the Crop diversification requirement and information about their crop fields

| Farmer id | crop types count | sum area of parcels | main crop(s) area | crop types |
|-----------|------------------|---------------------|-------------------|---|
| 37 | 1 | 245008 | | Soft Wheat |
| 1005 | 2 | 309458 | | Soft Wheat, Barley |
| 1007 | 1 | 191124 | | Soft Wheat |
| 1133 | 2 | 274066 | 271058 | Soft Wheat, Rape |
| 1229 | 2 | 407609 | | Soft Wheat, Barley |
| 1319 | 1 | 103471 | | Barley |
| 1306 | 1 | 302232 | | Soft Wheat |
| 1860 | 2 | 287242 | 278455 | Soft Wheat, Rape |
| 2677 | 3 | 577902 | 576071 | Rape, Soft Wheat, Vinification vineyard |
| 3131 | 3 | 614529 | 586915 | Soft Wheat, Rape, Barley |

Figure 10 presents the percentage of farmers that seem to be complying and non-complying to the Greening 1 requirement. The number of farmers that there is no need to be checked is 2610 (82%), while the number of potentially non-complying farmers is 703 (18%). More specifically, 506 (13%) farmers own 10 to 30 hectares of arable farm from which 54 (1%) grow less than two different crop types and 452 (14%) grow two or more different crop types but the main crop covers more than 75% of the land. On the other hand, 189 (5%) farmers own more than 30 hectares of arable farm from which 68 (2%) grow less than three different crop types and 121 (3%) grow three or more different crop types but the two main crops cover more than 95% of the land. To perform these calculations we used the data of the classification run from early July (jul08).




Fig. 10: Farmers distribution under Greening 1 requirement according to their compliance and total arable farm area that have under their possession

3) *Implementation of Scenario 3: SMR 1:* In the SMR 1 requirement we select the parcels that their distances from surface is under 10 meters based on the organic manure spreading buffer. From these only parcels with slope higher than 12° are recorded, according to the SMR1 specifications. Additionally, in order to select only parcels with an actual risk for runoff, the aspect of the proximity line ($\alpha_{degrees}$), namely the orientation of the line connecting the parcel to the water object, needs to fulfill the relationship:

$$aspect - 45 < \alpha_{degrees} < aspect + 45 \quad (6)$$

where the aspect is given per parcel and the proximity line aspect is calculated using the following formula:

$$\alpha_{radians} = \text{atan2}(x_1 - x_2, y_1 - y_2) \quad (7)$$

$$\alpha_{degrees} = \alpha_{radians} \times 180^\circ / \pi \quad (8)$$

Figure 11 shows the visual representation of the parcel that was mentioned in section IV-B3. According to the aforementioned rules this parcel is of high risk.

Figure 12 presents the percentage of parcels that seem to be of low and high risk based on the SMR 1 requirement. The number of low risk parcels that there is no need to be checked is 14407 (96%), while the number of high risk parcels is 630 (4%).

D. Smart Sampling Accuracy

Table V presents the PA and UA of the misclassified green parcels, for varying and constant (5) persistence, with the latter functioning as the ground truth. To calculate the PA, the number of correctly misclassified green parcels of varying persistence (P_i) is divided by the total number of misclassified parcels in the validation dataset (267), while for the UA it is divided by the total number of misclassified green parcels of varying persistence.

Table V shows that suboptimal PAs and UAs are achieved for classifications early in the year. This indicates that the smart sampling algorithm will select erroneously a significant percentage of the suggested OTSCs. Nevertheless, the smart



Fig. 11: Example of parcel that is susceptible to runoff according to the SMR 1 requirement. The parcel slope aspect (arrow) and the location of the nearby Regata de Larrea river are also shown.




Fig. 12: Distribution of low risk and high risk parcels according to the SMR1 specifications

TABLE V PA and UA evolution of the smart sampling algorithm

| Date | PA (%) | UA (%) |
|-------|--------|--------|
| May04 | 56.18 | 32.82 |
| May19 | 50.19 | 39.88 |
| May24 | 52.81 | 39.50 |
| Jun03 | 48.69 | 50.98 |
| Jun23 | 50.94 | 70.10 |
| Jul08 | 53.93 | 97.96 |
| Jul23 | 56.93 | 100.00 |
| Aug02 | 54.68 | 100.00 |
| Aug27 | 71.54 | 98.96 |
| Sep11 | 69.29 | 100.00 |
| Sep16 | 79.40 | 100.00 |
| Oct16 | 83.15 | 100.00 |
| Oct21 | 85.77 | 100.00 |

sampling precision, indicated by the UA, reaches near perfect values, from early July08 onward.

In an attempt to further filter the OTSCs selection of the suboptimal smart sampling results, found early in the year, we exploit the season type of the crop taxonomy, as described in Figure 5. Each of the crops that were examined belongs to one of the three season types and that is summer, winter, and year-round. Based on that, a false declaration would be more likely, if the prediction of the algorithm belongs to a different season class as compared to the corresponding season class of the declaration. Table VI presents the number of the green misclassifications that come from each of the different season types as well as the total sum of them, along with the number of green misclassifications for which the predicted crop season differs from the declared one, identified as an alarm. Finally, the last column indicates the percentage of those alarms that belong to the validation dataset of more than five persistent misclassifications.

TABLE VI Percentages of green misclassifications for each season class and UA of persistent alarms filtered through Algorithm

| Date | Total | Winter | Summer | All Year | Alarms | Persistent Alarms | UA (%) |
|-------|-------|--------|--------|----------|--------|-------------------|--------|
| May04 | 457 | 418 | 3 | 1 | 35 | 57.14 | |
| May19 | 336 | 302 | 3 | 1 | 30 | 63.33 | |
| May24 | 357 | 321 | 4 | 2 | 30 | 60.00 | |
| Jun03 | 255 | 226 | 4 | 1 | 24 | 70.83 | |
| Jun23 | 194 | 167 | 4 | 1 | 22 | 68.18 | |
| Jul08 | 147 | 125 | 3 | 1 | 18 | 94.44 | |
| Jul23 | 152 | 133 | 2 | 1 | 16 | 100.00 | |
| Aug02 | 146 | 128 | 1 | 1 | 16 | 100.00 | |
| Aug27 | 193 | 168 | 0 | 2 | 23 | 91.30 | |
| Sep11 | 185 | 164 | 0 | 2 | 19 | 100.00 | |
| Sep16 | 212 | 182 | 0 | 2 | 28 | 100.00 | |
| Oct16 | 222 | 186 | 1 | 2 | 33 | 100.00 | |
| Oct21 | 229 | 190 | 1 | 2 | 36 | 100.00 | |

Comparing the “Persistent Alarms UA” column of Table VI and the “UA” column of Table III reveals increase in the precision of smart sampling for runs that take place as early as May. Special attention is given to UA instead of PA, as it demonstrates the reliability of the system. Since, the inspections are not exhaustive but rather sampled based, it is more important to ensure that most alerts are indeed wrong declarations. In Figure 13 is displayed the evolution of the smart sampling alerts for different classifications throughout the year. It can be observed that for runs until June 23, the alerts are only few as they are passed through the Algorithm 2 filter. It can be seen that even though some alerts only appear early in the year, the critical mass of them can be identified from as early as July.

E. Scalability

Table VII presents the execution time for each one of rules that are described in this work. The execution time for each rule is calculated as the mean time of each scenario of the queries presented in section III-D. Results show that the query that takes the most time to run is the SMR 1 requirement, taking into account that this query runs per parcel, compared to the other queries which run for the whole dataset. The reason is that this query requires multiple calculations between polygons such as distance and aspect which are very time-consuming. The second column presents the execution time in seconds

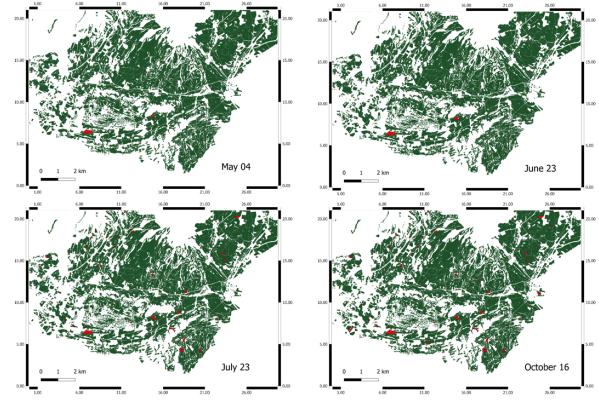


Fig. 13: Smart sampling alerts at different instances throughout the year of inspection. The red parcels indicate to the paying agency inspector where to target their inspections. Early in the year the alerts are fewer, stemming from less reliable classifications. With more images as we move along in time, classifications become more reliable and thus more alerts are identified.

for the actual size of the dataset, while the last presents the execution time in a reduced dataset size in order to further understand the scalability.

TABLE VII Mean execution time of each scenario according to the semantic queries described in section III-D

| Query | Execution time (sec) | Execution time (sec) |
|------------------------|----------------------|----------------------|
| | actual size | reduced size |
| Smart Sampling | 0.4 | 0.4 |
| Greening 1 requirement | 7.35 | 0.023 |
| SMR 1 requirement | 1.16 | 0.016 |
| Number of triples | 7.369.055 | 1.260.600 |




Fig. 14: SMR 1 requirement mean execution time for each parcel according to the risk type (low, high) compared with total mean execution time

Figure 14 presents the mean execution time of SMR 1 requirement for each parcel according to the risk type (low, high). The execution time is extremely low when the parcel does not satisfy the filters (aspect, distance, slope), while when the parcel satisfies the filters the execution time is significantly higher. Despite this fact, the mean execution time of all parcels

is significantly low, taking into account the large number of low-risk parcels.

F. CAP monitoring in practice

The results in Section IV-C described the implementation of the practical applications of the proposed system, which have been showcased in the form of the three scenarios of Section IV-B. Scenario 2 on Greening 1, and Scenario 3 on SMR-1 requirements respectively, are performed once for every year of inspection. Therefore, these two scenarios infer a single cycle of execution, with reference to Figure 1. On the other hand, the Smart Sampling application (Scenario 1) is an iterative process. The first crop classification is performed in early May, coinciding with the CAP subsidies applications commencement, with subsequent iterations producing progressively more accurate crop type maps with every new Sentinel-2 acquisition. Hence, there is a full cycle execution (Figure 1) with every newly acquired image, from ingestion to interlinking.

According to the results in Table VI, the issued alarms are adequately trustworthy to suggest potential OTSCs from as early as July, with a user's accuracy of more than 94%. From then on, the paying agency inspectors can have targeted OTSCs that increase in number with every new iteration, following their summer-long inspection process. Alternatively, the results from May until early July can be used in assistance of the farmer application process. The applicants and the paying agencies can have an indication of potential non-compliance even at the application stage; thus allowing the farmer to make timely changes to their application.

V. CONCLUSION AND FUTURE WORK

In this paper we presented a semantic-oriented framework for knowledge discovery using supervised classification in the CAP domain. The main focus is given on the detection of possible violations according to the declaration of the farmers, the Greening 1 requirement and SMR 1 requirement. The framework can strongly assist in decision making issues by providing helpful information to paying agency inspectors and environmental consultant to detect possible breaches of compliance. The proposed solution relies on data coming from sentinel images and open data (e.g. OpenStreetMap). Common denominator between the two datasets is the provision of georeferenced information. Data refer to a region in northeastern Spain. The SVM classification method has been applied to classify the cultivated crop types for multiple instances throughout the cultivation season. The data, which are in Shapefile format, are converted into Turtle RDF format using the GeoTriples tool. Results are saved in GraphDB triplestore. Semantic queries are executed to enrich the data with information about possible farmers non-compliance according to agricultural policies.

In this work we have shown how the paying agencies of the CAP can benefit from the exploitation of big Copernicus data. We showcased how with only freely available satellite data and ancillary LOD one can provide actionable information. Combining the state-of-the-art in EO-based crop classification,

semantic enrichment and linking free and open data has facilitated the development of an end-to-end system, from data acquisition to CAP related decision making. The main innovations of the presented methodology include its re-usability and transferability, using predominantly open data and requiring minimal fine-tuning when applied to other regions, and scalability, accounting for all big data considerations and choosing the computationally efficient alternative every step of the way, towards the monitoring approach of the new CAP.

Future work includes the investigation of similar datasets that correspond to other regions to apply the agricultural policies rules. The proposed framework can also be extended to similar problems, beyond the control of the EU Common Agricultural Policy, using semantic enrichment and reasoning to support farmers in monitoring their crops, insurance companies to assess the risk in a specific area of interest, and public agencies that monitor the sustainability of rural areas as a consequence of climate change.

ACKNOWLEDGMENT

This work has been supported by EOPEN, aqua3S and RECAP projects, which have been funded from the European Union's Horizon 2020 research and innovation programme under grant agreements 776019, 832876 and 693171 respectively. We are also grateful to INTIA for supplying the farmers' declarations, LPIS data and phenology timeline.

REFERENCES

- [1] Brian McBride. The resource description framework (rdf) and its vocabulary description language rdfls. In *Handbook on ontologies*, pages 51–65. Springer, 2004.
- [2] Niklas Kolbe, Sylvain Kubler, Jérémy Robert, Yves Le Traon, and Arkady Zaslavsky. Linked vocabulary recommendation tools for internet of things: a survey. *ACM Computing Surveys (CSUR)*, 51(6):1–31, 2019.
- [3] Manolis Koubarakis, Kostis Kyzirakos, Charalampos Nikolaou, George Garbis, Konstantina Bereta, Roi Dogani, Stella Giannakopoulou, Panayiotis Smeros, Dimitrianos Savva, George Stamoulis, et al. Managing big, linked, and open earth-observation data: Using the teleios/leo software stack. *IEEE Geoscience and Remote Sensing Magazine*, 4(3):23–37, 2016.
- [4] Yunqiang Zhu, A-Xing Zhu, Jia Song, Jie Yang, Min Feng, Kai Sun, Jingqu Zhang, Zhiwei Hou, and Hongwei Zhao. Multidimensional and quantitative interlinking approach for linked geospatial data. *International Journal of Digital Earth*, 10(9):923–943, 2017.
- [5] Manolis Koubarakis, Konstantina Bereta, George Papadakis, Dimitrianos Savva, and George Stamoulis. Big, linked geospatial data and its applications in earth observation. *IEEE Internet Computing*, 21(4):87–91, 2017.
- [6] Guy Pe'Er, Yves Zinngrebe, Francisco Moreira, Clélia Sirami, Stefan Schindler, Robert Müller, Vasileios Bontzorlos, Dagmar Clough, Peter Bezák, Aletta Bonn, et al. A greener path for the eu common agricultural policy. *Science*, 365(6452):449–451, 2019.
- [7] Valentine Lebourgeois, Stéphane Dupuy, Élodie Vintrou, Maël Ameline, Suzanne Butler, and Agnès Bégué. A combined random forest and obia classification scheme for mapping smallholder agriculture at different nomenclature levels using multisource data (simulated sentinel-2 time series, vhrs and dem). *Remote Sensing*, 9(3):259, 2017.
- [8] Vasileios Sitokoustantinou, Ioannis Papoutsis, Charalampos Kontoes, Alberto Lafarga Arnal, Armesto Andrés, Ana Pilar, and José Angel Garraza Zurbano. Scalable parcel-based crop identification scheme using sentinel-2 data time-series for the monitoring of the common agricultural policy. *Remote Sensing*, 10(6):911, 2018.
- [9] James Brinkhoff, Justin Vardanega, and Andrew J Robson. Land cover classification of nine perennial crops using sentinel-1 and-2 data. *Remote Sensing*, 12(1):96, 2020.

- [10] María Arias, MA Campo-Bescós, and Jesús Álvarez-Mozos. Crop type mapping based on sentinel-1 backscatter time series. In *IGARSS 2018-2018 IEEE International Geoscience and Remote Sensing Symposium*, pages 6623–6626. IEEE, 2018.
- [11] Siwen Feng, Jianjun Zhao, Tingting Liu, Hongyan Zhang, Zhengxiang Zhang, and Xiaoyi Guo. Crop type identification and mapping using machine learning algorithms and sentinel-2 time series data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(9):3295–3306, 2019.
- [12] Elahe Akbari, Ali Darvishi Boloorani, Najmeh Neysani Samany, Saeid Hamzeh, Saeid Soufizadeh, and Stefano Pignatti. Crop mapping using random forest and particle swarm optimization based on multi-temporal sentinel-2. *Remote Sensing*, 12(9):1449, 2020.
- [13] Yady Tatiana Solano-Correa, Francesca Bovolo, and Lorenzo Bruzzone. A semi-supervised crop-type classification based on sentinel-2 ndvi satellite image time series and phenological parameters. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 457–460. IEEE, 2019.
- [14] SS Ray et al. Exploring machine learning classification algorithms for crop classification using sentinel 2 data. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*, 2019.
- [15] Jitendra Singh, Aniruddha Mahapatra, Saurav Basu, and Biplab Banerjee. Assessment of sentinel-1 and sentinel-2 satellite imagery for crop classification in indian region during kharif and rabi crop cycles. In *IGARSS 2019-2019 IEEE International Geoscience and Remote Sensing Symposium*, pages 3720–3723. IEEE, 2019.
- [16] Vittorio Mazzia, Aleem Khaliq, and Marcello Chiaberge. Improvement in land cover and crop classification based on temporal features learning from sentinel-2 data using recurrent-convolutional neural network (r-cnn). *Applied Sciences*, 10(1):238, 2020.
- [17] Jingcheng Zhang, Yuhang He, Lin Yuan, Peng Liu, Xianfeng Zhou, and Yanbo Huang. Machine learning-based spectral library for crop classification and status monitoring. *Agronomy*, 9(9):496, 2019.
- [18] Deborah L McGuinness, Frank Van Harmelen, et al. Owl web ontology language overview. *W3C recommendation*, 10(10):2004, 2004.
- [19] Peter Patel-Schneider and Jérôme Siméon. The yin/yang web: Xml syntax and rdf semantics. In *Proceedings of the 11th international conference on World Wide Web*, pages 443–453, 2002.
- [20] Christian Bizer, Tom Heath, and Tim Berners-Lee. Linked data: The story so far. In *Semantic services, interoperability and web applications: emerging concepts*, pages 205–227. IGI Global, 2011.
- [21] Jorge Pérez, Marcelo Arenas, and Claudio Gutierrez. Semantics and complexity of sparql. *ACM Transactions on Database Systems (TODS)*, 34(3):1–45, 2009.
- [22] Manolis Koubarakis and Kostis Kyzirakos. Modeling and querying metadata in the semantic sensor web: The model strdf and the query language stsparql. In *Extended Semantic Web Conference*, pages 425–439. Springer, 2010.
- [23] Kostis Kyzirakos, Dimitrios Savva, Ioannis Vlachopoulos, Alexandros Vasileiou, Nikolaos Karalis, Manolis Koubarakis, and Stefan Manegold. Geotriples: Transforming geospatial data into rdf graphs using r2rml and rml mappings. *Journal of Web Semantics*, 52:16–32, 2018.
- [24] Altti Ilari Maarala, Xiang Su, and Jukka Riekk. Semantic reasoning for context-aware internet of things applications. *IEEE Internet of Things Journal*, 4(2):461–473, 2016.
- [25] Franz Baader, Ian Horrocks, and Ulrike Sattler. Description logics. *Foundations of Artificial Intelligence*, 3:135–179, 2008.
- [26] Helle V Dam, Jan Engberg, and Heidrun Gerzymisch-Arbogast. *Knowledge systems and translation*, volume 7. Walter de Gruyter, 2011.
- [27] Chao Wang, Wei Wang, and Nengcheng Chen. Building an ontology for hydrologic monitoring. In *2017 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 6232–6234. IEEE, 2017.
- [28] Maroua Masmoudi, Hela Taktak, Sana Ben Abdallah Ben Lamine, Mohamed Hedi Karray, Hajer Baazaoui Zghal, Bernard Archimede, Michael Mrissa, Chirine Ghedira Guegan, et al. Predicat: a semantic service-oriented platform for data interoperability and linking in earth observation and disaster prediction. In *2018 IEEE 11th Conference on Service-Oriented Computing and Applications (SOCA)*, pages 194–201. IEEE, 2018.
- [29] Gelian Song, Maohua Wang, Xiao Ying, Rui Yang, and Binyun Zhang. Study on precision agriculture knowledge presentation with ontology. *AASRI Procedia*, 3:732–738, 2012.
- [30] Michael T Maliappis. Applying an agricultural ontology to web-based applications. *International Journal of Metadata, Semantics and Ontologies*, 4(1-2):133–140, 2009.
- [31] Siquan Hu, Haiou Wang, Chundong She, and Junfeng Wang. Agont: ontology for agriculture internet of things. In *International Conference on Computer and Computing Technologies in Agriculture*, pages 131–137. Springer, 2010.
- [32] Ya-ru SU, Ru-jing WANG, CHEN Peng, Yuan-yuan WEI, Chuan-xi LI, and Yi-min HU. Agricultural ontology based feature optimization for agricultural text clustering. *Journal of Integrative Agriculture*, 11(5):752–759, 2012.
- [33] Surya S Durbha and Roger L King. Semantics-enabled framework for knowledge discovery from earth observation data archives. *IEEE Transactions on geoscience and remote sensing*, 43(11):2563–2572, 2005.
- [34] Garima Jaiswal and Arun Prakash Agrawal. Comparative analysis of relational and graph databases. *IOSR Journal of Engineering (IOSRJEN)*, 3(8):2250–3021, 2013.
- [35] Georgios Meditskos and Nick Bassiliades. Rule-based owl ontology reasoning systems: Implementations, strengths, and weaknesses. In *Handbook of Research on Emerging Rule-Based Languages and Technologies: Open Solutions and Approaches*, pages 124–148. IGI Global, 2009.
- [36] Kostis Kyzirakos, Manos Karpathiotakis, and Manolis Koubarakis. Strabon: a semantic geospatial dbms. In *International Semantic Web Conference*, pages 295–311. Springer, 2012.
- [37] Ralf Hartmut Güting. Graphdb: Modeling and querying graphs in databases. In *VLDB*, volume 94, pages 12–15. Citeseer, 1994.
- [38] Pierfrancesco Bellini and Paolo Nesi. Performance assessment of rdf graph databases for smart city services. *Journal of Visual Languages & Computing*, 45:24–38, 2018.
- [39] Orri Erling and Ivan Mikhailov. Rdf support in the virtuoso dbms. In *Networked Knowledge-Networked Media*, pages 7–24. Springer, 2009.
- [40] Konstantina Bereta, Guohui Xiao, and Manolis Koubarakis. Onto-spatial: Ontop of geospatial databases. *Journal of Web Semantics*, 58:100514, 2019.
- [41] Jitao Yang and Guoqing Li. Earth observation data integration and opening system. In *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 5481–5484. IEEE, 2016.
- [42] Efstratios Kontopoulos, Panagiotis Mitziias, Jürgen Moßgraber, Philipp Hertweck, Hylke van der Schaaf, Désirée Hilbring, Francesca Lombardo, Daniele Norbiato, Michele Ferri, Anastasios Karakostas, et al. Ontology-based representation of crisis management procedures for climate events. In *ISCRAM*, 2018.
- [43] Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Joint learning from earth observation and openstreetmap data to get faster better semantic maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 67–75, 2017.
- [44] Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Fusion of heterogeneous data in convolutional networks for urban semantic labeling. In *2017 Joint Urban Remote Sensing Event (JURSE)*, pages 1–4. IEEE, 2017.
- [45] Wei Yao, Dimitrios Marmanis, and Mihai Datcu. Semantic segmentation using deep neural networks for sar and optical image pairs. 2017.
- [46] Désirée Hilbring, Jürgen Moßgraber, Philipp Hertweck, Tobias Hellmund, et al. Harmonizing data collection in an ontology for a risk management platform.
- [47] Victor Francisco Rodríguez-Galiano, Bardan Ghimire, John Rogan, Mario Chica-Olmo, and Juan Pedro Rigol-Sanchez. An assessment of the effectiveness of a random forest classifier for land-cover classification. *ISPRS Journal of Photogrammetry and Remote Sensing*, 67:93–104, 2012.
- [48] Bo-Cai Gao. NdwI—a normalized difference water index for remote sensing of vegetation liquid water from space. *Remote sensing of environment*, 58(3):257–266, 1996.
- [49] Ting-Fan Wu, Chih-Jen Lin, and Ruby C Weng. Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research*, 5(Aug):975–1005, 2004.
- [50] Marina Sokolova, Nathalie Japkowicz, and Stan Szpakowicz. Beyond accuracy, f-score and roc: a family of discriminant measures for performance evaluation. In *Australasian joint conference on artificial intelligence*, pages 1015–1021. Springer, 2006.