Fire Risk Management using Data Cubes, Machine Learning and OBDA systems (Demo Paper)

Dimitris Bilidas¹, Anastasios Mantas¹, Filippos Yfantis¹, George Stamoulis¹, Manolis Koubarakis¹, Spyros Kondylatos², Ioannis Prapas², Ioannis Papoutsis²

¹Dept. of Informatics and Telecommunications, National and Kapodistrian University of Athens, Greece {d.bilidas, cs2190004, cs2190003, gstam, koubarak}@di.uoa.gr ²Orion Lab, National Observatory of Athens, Greece {skondylatos, iprapas, ipapoutsis}@noa.gr

ABSTRACT

We present a fire risk management system which takes input data from various sources (e.g., meteorological data, satellite indicators for vegetation, historical burned areas), produces a harmonized spatio-temporal data cube to compute fire risk and enables semantic querying to assist fire risk management. The distinguishing implementation features of the system is the use of data cubes, machine learning algorithms and, most importantly, geospatial ontology-based data access technologies. The system has been implemented in the European project DeepCube for the geographic area of Greece and can be used operationally to assist authorities to determine fire risk during the summer fire season.

CCS CONCEPTS

• Information systems → Information systems applications; Spatial-temporal systems; Geographic information systems;

KEYWORDS

data cubes, ontology based data access, machine learning, fire risk

ACM Reference Format:

Dimitris Bilidas¹, Anastasios Mantas¹, Filippos Yfantis¹, George Stamoulis¹, Manolis Koubarakis¹, Spyros Kondylatos², Ioannis Prapas², Ioannis Papoutsis² . 2023. Fire Risk Management using Data Cubes, Machine Learning and OBDA systems (Demo Paper). In *The 31st ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '23), November 13–16, 2023, Hamburg, Germany.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3589132.3625615

1 INTRODUCTION

Wildfires pose a significant global natural hazard, disturbing natural ecosystems, leading to loss of life, property, and infrastructure, while also contributing to the emission of carbon dioxide. Climate change is progressively influencing wildfire patterns and is projected to intensify wildfires across a majority of the globe, expanding the risk of wildfires to higher latitudes, evergreen tropical forests, and notably in the broader Mediterranean region. Particularly, scenarios for global warming greater than 1.5°C could lead to a 40% increase in Mediterranean burned area [14]. To address

SIGSPATIAL '23, November 13–16, 2023, Hamburg, Germany © 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0168-9/23/11.

https://doi.org/10.1145/3589132.3625615

the new challenges, it is important to make use of multifaceted sources of data and develop new technologies for timely and accurate fire hazard forecasting. In this demo paper, we present a fire risk management system that utilizes data cubes, machine learning algorithms and geospatial ontology-based data access (OBDA) technologies to compute fire risk.

A *data cube* is a multidimensional array of values, serving as a natural data structure for storing Earth observation (EO) data and various other forms of multidimensional data for analysis purposes. Several data cube infrastructures have emerged specifically targeting EO data, such as the Open Data Cube infrastructure in Australia, and the Euro Data Cube and the Earth System Data Cube funded by the European Space Agency. These infrastructures provide libraries and APIs, such as xarray and YAXArrays, designed for efficient storage and querying of multidimensional data. However, prior to the widespread adoption of data cube infrastructures, extensive research and development had already taken place on array database management systems (DBMS) like Rasdaman [2], SciDB [13], and MonetDB SciQL [17]. These DBMS, in contrast to data cube infrastructures, offer *declarative query languages* tailored to model and query multidimensional data.

The system Plato used in the fire risk management system of this paper goes beyond data cube infrastructures and array database management systems, and it is the first semantic data cube system implemented using geospatial ontology-based data access technologies [3]. The idea of semantic data cube systems was initially introduced in [1] but, in our opinion, it did not receive the attention it deserved. The term semantic was employed in [1] to distinguish semantic data cubes from regular ones constructed using the aforementioned infrastructures and array DBMS, which primarily contain numerical values (e.g., 10-day average land surface temperature above 35°C). In semantic data cubes, these values are associated with symbolic high-level concepts (e.g., a heat wave). Beyond providing knowledge through interpretations (as in the example just given), semantic data cube systems facilitate the incorporation of external datasets, enabling combined analyses. For instance, demographic data published by a government organization can be utilized to identify major cities located within a certain distance from areas where the risk of a wildfire is significant during a particular time of the year.

Plato is a pioneering semantic data cube system utilizing geospatial OBDA technologies. OBDA serves as a methodology for connecting an ontology, which captures geospatial knowledge about entity classes and properties within a specific application domain, to underlying data sources. These data sources, managed by specialized systems, exist in various formats and typically reside in

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

pre-existing repositories (e.g., a geospatial relational DBMS or a shapefile). To establish the connection, declarative mappings are employed, enabling the generation of ontology terms based on information extracted from the data sources. Rather than materializing all ontology terms, in OBDA systems users can pose queries directly on the ontology. Subsequently, a process of query transformation takes place, converting the user's query into the native language understood by the underlying data sources (e.g., GeoSPARQL). This transformed query is executed, and the results are then converted back into ontology terms for presentation to the user. This approach, often referred to as the *virtual knowledge graph* approach, offers the advantage of providing users with a familiar vocabulary to articulate queries while abstracting the complexities of the underlying data sources, including intricate schemas and storage nuances.

Plato is implemented using the OBDA system Ontop [16], one of the pioneer OBDA systems capable of performing SPARQL to SQL query translation. The process involves several inputs: (i) an ontology expressed in the OWL2 QL subset of the OWL2 ontology language [15], (ii) a database schema, (iii) a set of mapping assertions that generate virtual RDF triples from database values, and (iv) an initial SPARQL query targeting the ontology. As a result, an SQL query is generated, which can be executed on any database instance conforming to the input schema, providing comprehensive answers consistent with the ontology axioms.

In 2016, the University of Athens team leading this paper introduced Ontop-spatial [3, 4] as the first *geospatial* OBDA system. It was developed as a geospatial extension of Ontop. Within Ontopspatial, a GeoSPARQL query is transformed into an intermediate representation based on Datalog. This query is then rewritten, taking into account the ontology and mappings from ontology concepts to data sources. The final outcome is an SQL query that utilizes spatial SQL functions, corresponding to the GeoSPARQL functions and operators in the initial query. This SQL query can be executed within a spatially-enabled relational system, such as PostGIS (the spatial extension of PostgreSQL) or Spatial-Lite (the spatial extension of SQLite). Starting from version 4.1.0, the functionalities of Ontop-spatial have been fully integrated into Ontop.

2 SYSTEM ARCHITECTURE

In this section we present the different components of the fire risk management system as it has been developed in the context of European project DeepCube (https://deepcube-h2020.eu/). The architecture of the system is shown in Figure 1.



Figure 1: System architecture

The system takes as input several sources of satellite and vector data, as shown in Table 1, to produce a data cube that is used to calculate fire risk using machine learning (ML) models. Since the input consists of satellite images/products, the resulting cube contains cells with geospatial coordinates, corresponding to a pixel from the input image, along with their attributes. This allows us to calculate fire risk for specific areas and visualize the results as interactive maps. Fire risk information is then combined with other vector sources (e.g., Natura areas, OpenStreetMap) and allows the user to pose semantic queries using the query language GeoSPARQL, with the use of Plato.

The data cube, named FireCube [11], is populated with data collected for the area of Greece, for the years 2009-2021. It contains 90 variables related to the drivers of fire occurrence and spread. Weather data are extracted from the dataset ERA-5 Land¹, satellite variables as proxies of vegetation status and dynamics are extracted from the Moderate Resolution Imaging Spectroradiometer (MODIS) aboard the Terra and Aqua satellites, and human activity proxies are collected from the WorldPop dataset². The data cube is annotated with the historical burned areas³ from the European Forest Fire Information System (EFFIS) product and further augmented with the EFFIS active fires product³ in order to extract the actual wildfire ignition date. Finally, all the data is harmonized into a *1km x 1km x 1day* spatio-temporal data cube, which covers the whole of Greece (*1253km x 983km*); see [8] for more details.

Training Data. From the above data cube we extract four different ent datasets [8] that are used to train four different models presented below. For a given pixel (i.e., a cell representing a *1km x 1km* square region) and a given day, we extract input-target pairs for 4 different modeling modalities:

- A pixel dataset, where we extract the input attributes and their last 10-day average.
- A temporal dataset, where we extract the last 10-day timeseries of the input attributes.
- A spatial dataset, where we extract *25km x 25km* patches spatially centered around the given pixel.
- A spatio-temporal dataset, where we extract 25km x 25km x 10days blocks centered spatially around the given pixel.

Machine Learning Models. The complexity of the interactions of the input variables and the stochastic nature of wildfire behaviour motivate us to use machine learning methods which are known for their ability to leverage this complexity and to identify information hidden in the data. These ML methods can assist in wildfire danger prediction for assessing the conditions that allow a fire to ignite and spread. We train a different model for each type of the abovementioned machine learning datasets.

For the pixel dataset, we consider a Random Forest baseline model. Its hyperparameters are chosen based on the validation Area Under the Receiver Operating Characteristic (AUROC). For the temporal dataset, we use a Long Short-Term Memory (LSTM) architecture that is able to capture temporal dynamics. For the spatial dataset, we use a CNN architecture, known to effectively capture spatial features. For the spatio-temporal dataset, we use a

 $[\]label{eq:linear} ^1 https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-land?tab=form ^2 https://hub.worldpop.org/project/categories?id=18$

³https://effis.jrc.ec.europa.eu/applications/data-and-services

Dataset	Description	Spatial Resolution	Temporal Resolution
MOD11	Land Surface		
	Temperature and	1km	daily
	Emissivity (Land Surface		
	Temperature Day &		
	Night)		
MOD13	Vegetation Index	1km	16-day
	Products (Normalized		
	Difference Vegetation		
	Index, Enhanced		
	Vegetation Index)		
MOD15	Radiation, Leaf Area	500m	8-day
	Index		
MOD16	Evapotranspiration	500m	8-day
ERA5-Land	Wind speed & direction,	9km	5 hours/ daily
	precipitation,		
	temperature, surface		
	pressure, dew point		
	temperature		
EDO	Soil Moisture Index, Soil	5km	10-day
	Moisture Anomaly		
CLC	Corine Land Cover	100m	2006/2012/
			2018
EU-DEM	Digital Elevation Model	25m	-
Aspect	Aspect	25m	-
Slope	Slope	25m	-
WorldPop	Population Density	1km	2009-2020
OSM	Road Distance,	vector	-
	Waterway Distance		
EFFIS	Burned Areas provided by EFFIS	vector	2009-2021
Burned			
Areas			
MODIS			
Active Fire	Fire Hotspots	points	2009-2021
Data			

Table 1: Data cube input sources.

ConvLSTM, which is known to capture spatiotemporal features of the input, combining the strengths of the LSTM and CNN components. We compare the predictive skill of these four models over two distinct test sets for years 2020 and 2021, vis-à-vis the baseline of the empirical Fire Weather Index which relies solely on meteorological conditions and disregards the status of other fire drivers related to vegetation and human factors. Our experimental analysis shows that the ConvLSTM model that considers the spatiotemproal context leading to a wildfire performs best in forecasting fire danger [8].

The Semantic Data Cube System Plato. Plato consists of two main components, the OBDA system Ontop and a PostGIS backend. To initialize the Ontop engine, we provide an ontology in the OWL2 ontology language, along with a specified set of mappings. Ontologies provide a familiar vocabulary for the user in terms of classes and properties, while mappings define the way the ontology terms are related to the data residing in the backend. After initialization, Ontop is ready to accept GeoSPARQL queries and translate them into SQL enhanced with spatial operators. The PostGIS backend contains the virtual tables that are used to communicate with local or remote data cubes. To implement this communication, we utilize foreign data wrappers (FDWs) that are based on the Xarray [7] library and the Multicorn package [6]. To handle large volumes of data and the heavy computations in spatial joins between raster and vector data, Plato implements two optimization techniques: data caching and Raptor Join [12].

The idea behind caching raster data is to optimize the retrieval of large data cubes by pre-computing and storing relevant portions. This approach ensures that the required data is readily accessible, thereby enhancing query efficiency. During the translation of queries from GeoSPARQL to SQL, Plato identifies the specific sections of raster data that need to be accessed and converted into geometries. These identified portions are then stored in an intermediate cache table within the database. This streamlined access to "hot" data is the key advantage of this implementation, enabling more efficient joins and other operations between dense data cubes and other materialized (non-EO) data.

The second optimization of Plato is the use of Raptor Join [12] to overcome the bottleneck of joining large amounts of raster data that are part of the data cube and vector data that need to be combined with the cube for the purposes of our application. Raptor Join does this by selectively reading only the parts of the raster that intersect with a set of vector geometries using *scanlines*. This method eliminates the need for conversions between raster and vector representations for join operations. In Plato, we have implemented the Raptor Join as a Python FDW that computes the output of spatial operations.

The implementation of Plato and its performance evaluation is presented in [5].

3 PRESENTATION AND DEMONSTRATION

Our presentation will first introduce the main ideas of the fire risk management system as they have been discussed in this paper. Then, our demonstration will focus on the use of Plato and showcase how GeoSPARQL queries can be used to combine information from data cubes and vector data to provide valuable insight for authorities (e.g., the Greek fire brigade) so that they can assess the risk of fire and act accordingly. These queries essentially formalize "golden rules" that operational fire risk managers have developed over the years and use them to take decisions regarding fire risk. The queries to be demonstrated combine fire risk information for a specific area, as it has been computed by the ML models and stored in the data cube, with other sources (such as protected areas, POIs, nearby water sources, road network, land use/cover, etc.). Four such queries are given below:

- Q1. Which Natura areas are surrounded by mountainous areas (use elevation to find mountainous) and experience extreme fire danger (>0.9 risk) for a specific time period?
- Q2. Find OpenStreetMap POIs that are less than a specified distance away from areas with fire prediction greater than a given value.
- Q3. Which forests of type *X* that have already been burnt at least once in the past 10 years, are in high danger tomorrow?
- Q4. Which areas (or assets: population, industrial zones, natura areas, etc.) are in high danger due to forecasted strong winds and which due to dry conditions (low relative humidiy, low soi moisture, etc.) that are persistent for the last 10 days?

To formulate the queries in Plato, we have designed an ontology⁴ and made it publicly available on zenodo. During the demonstration, the ontology will be shown to the users using Protégé [9]. The three main classes in the ontology are Observation, Prediction and SpatialObject. The subclasses of class Observation (Daily Land Surface Temperature, Daily Surface Temperature, etc.) correspond to variables of the data cube. The subclass of class Prediction, FirePred corresponds to the daily fire predictions. Each observation and each prediction has a specific acquisition date and is related to a specific RasterCell, which is a subclass of SpatialObject. Each RasterCell has a point geometry corresponding to the latitude and longitude dimensions of the data cube. The data cube is accessed using an FDW in PostGIS. The subclasses of Feature (AdministrativeUnit, NaturaArea, etc.) correspond to information that is integrated from other data sources (such as shapefiles and TIFF images). These data sources can be imported in PostGIS. During the demonstration, the necessary datasets will be imported from 11 shapefiles and 10 TIFF files. These include two layers of the administrative units of Greece (regions and municipalities), OpenStreetMap data that correspond to particular classes of the ontology like Roads (road network, railways), Settlements (different buildings), Water (bodies of water, streams) and FireProneStructure (power stations, fuel stations), Natura protected areas of Greece represented by the ontology class NaturaArea, and a total of ten TIFF files with daytime and nighttime fire prediction indices for five days (two TIFF files for each day) that correspond to the ontology class FirePred.

In order to perform query answering, we will provide a file with all the mappings from the ontology to the underlying database for the above-mentioned sources. As an example, the mapping in Figure 2 is used to connect the RasterCell class to information coming from the data cube. Each mapping has a source SQL query, and a target triples template. The specific source query selects the latitude, longitude and time dimension for each cell in the data cube. The target triples template dictates how triples are (virtually) generated from the results of the source query. Mappings like the ones in Figure 2 will be shown using Protégé.



Figure 2: RasterCell Ontology Class mapping

Plato takes as input the ontology and mappings and provides a SPARQL endpoint to pose queries and retrieve the results. In our demonstration users will also be able to pose their own queries, utilizing the presented ontology and mappings, view the query translation step and check how data are accessed from the data cube through the FDW and combined with the imported data in PostGIS. In Figure 3 we showcase a map in the tool Sextant [10], that presents layers on the map with the results of the queries for this demo. Sextant will be used during the demonstration to visualize the results of GeoSPARQL queries such as the ones above.



Figure 3: Fire risk management map in Sextant

ACKNOWLEDGEMENT

This work was supported by the first call for H.F.R.I. Research Projects to support faculty members and researchers and the procurement of high-cost research equipment grant (HFRI-FM17-2351). It was also partially supported by the ESA project DA4DTE (subcontract 202320239), the Horizon 2020 projects AI4Copernicus (GA No. 101016798) and DeepCube (GA No. 101004188), and Horizon Europe project STELAR (GA No. 101070122).

REFERENCES

- H. Augustin, M. Sudmanns, D. Tiede, S. Lang, and A. Baraldi. 2019. Semantic Earth Observation Data Cubes. *Data* 4, 3 (2019).
- [2] P. Baumann, P. Furtado, R. Ritsch, and N. Widmann. 1997. The RasDaMan approach to multidimensional database management. In ACM SAC.
- [3] K. Bereta and M. Koubarakis. 2016. Ontop of Geospatial Databases. In ISWC.
 [4] K. Bereta, G. Xiao, and M. Koubarakis. 2019. Ontop-spatial: Ontop of geospatial
- databases. J. Web Semant. 58 (2019).
 [5] Dimitris Bilidas, Anastasios Mantas, Filippos Yfantis, George Stamoulis, and Manolis Koubarakis. [n. d.]. Plato: A Semantic Data Cube Implementation using
- Ontology-Based Data Access Technologies. Under review for a conference. [6] R. Dunklau and F. Mounier. 2015. Multicorn - PostgreSQL extension. https://
- //multicorn.org/
 [7] S. Hoyer and J. Hamman. 2017. xarray: N-D labeled arrays and datasets in Python. Journal of Open Research Software 5, 1 (2017).
- [8] Spyros Kondylatos, Ioannis Prapas, Michele Ronco, Ioannis Papoutsis, Gustau Camps-Valls, María Piles, Miguel-Ángel Fernández-Torres, and Nuno Carvalhais. 2022. Wildfire danger prediction and understanding with Deep Learning. *Geophysical Research Letters* 49, 17 (2022), e2022GL099368.
- [9] Mark A. Musen. 2015. The protégé project: a look back and a look forward. AI Matters 1, 4 (2015), 4–12. https://doi.org/10.1145/2757001.2757003
- [10] Charalampos Nikolaou, Kallirroi Dogani, Konstantina Bereta, George Garbis, Manos Karpathiotakis, Kostis Kyzirakos, and Manolis Koubarakis. 2015. Sextant: Visualizing time-evolving linked geospatial data. J. Web Semant. 35 (2015), 35–52. https://doi.org/10.1016/j.websem.2015.09.004
- [11] I Prapas, S Kondylatos, and I Papoutsis. 2022. FireCube: A daily datacube for the modeling and analysis of wildfires in Greece. https://doi.org/10.5281/zenodo. 4943354
- [12] S. Singla, A. Eldawy, T. Diao, A. Mukhopadhyay, and E. Scudiero. 2021. The Raptor Join Operator for Processing Big Raster + Vector Data. In ACM SIGSPATIAL.
- [13] Michael Stonebraker, Paul Brown, Alex Poliakov, and Suchi Raman. 2011. The Architecture of SciDB. In SSDBM.
- [14] Marco Turco, Juan José Rosa-Cánovas, Joaquín Bedia, Sonia Jerez, Juan Pedro Montávez, Maria Carmen Llasat, and Antonello Provenzale. 2018. Exacerbated fires in Mediterranean Europe due to anthropogenic warming projected with non-stationary climate-fire models. *Nature Communications* 9, 1 (December 2018), 1–9. https://doi.org/10.1038/s41467-018-06358-
- [15] W3C. 2012. OWL 2 Web Ontology Language Profiles (Second Edition). http: //www.w3.org/TR/owl2-profiles/
- [16] G. Xiao, D. Lanti, R. Kontchakov, S. Komla-Ebri, E. Güzel-Kalaycı, L. Ding, J. Corman, B. Cogrel, D. Calvanese, and E. Botoeva. 2020. The Virtual Knowledge Graph System Ontop. In *ISWC*.
- [17] Y. Zhang, M. L. Kersten, and S. Manegold. 2013. SciQL: array data processing inside an RDBMS. In ACM SIGMOD.

⁴https://zenodo.org/record/7657265/files/UC3-Ontology.owl?download=1