Giving meaning to unsupervised EO change detection rasters: a semantic-driven approach

Jordane Dorne
jordane.dorne@thalesalieniaspace.com
Thales Alenia Space
France and IRIT - CNRS, Toulouse
University
France

Nathalie Aussenac-Gilles
Catherine Comparot
Cassia Trojahn
nathalie.aussenac-gilles@irit.fr
catherine.comparot@irit.fr
cassia.trojahn@irit.fr
IRIT - CNRS, Toulouse University
France

Romain Hugues
romain.hugues@thalesalieniaspace.com
Thales Alenia Space
France

ABSTRACT

The field of Earth Observation (EO) change detection has been fostered with new sources of satellite image data coupled with the development of deep learning algorithms. However, the output of these algorithms lacks context. Contextual knowledge explaining a detected change is required to better analyze those images and understand the phenomenon that caused the change. This paper presents a semantic-driven data integration approach that supports the generation of a knowledge graph from a raster change file and from various data sources of events that may explain the changes. The output graph represents spatial and temporal features for areas affected by a high change, as well as various kinds of contextual data useful for explaining the detected changes. We validate the approach with a real-case scenario of fire monitoring. We process changes detected between pairs of Sentinel-2 images located on the same tiles, with contextual data such as administrative units, tweets, and thermal sensors. We show the added value of the proposed approach for i) explaining change detection and ii) validating the results from unsupervised deep learning algorithms.

CCS CONCEPTS

• **Information systems** → *Information integration*; • **Computing methodologies** → *Ontology engineering*; • **Applied computing** → *Earth and atmospheric sciences*; • **Theory of computation** → *Data integration*.

KEYWORDS

Satellite images, semantic integration, change detection, knowledge graphs.

ACM Reference Format:

Jordane Dorne, Nathalie Aussenac-Gilles, Catherine Comparot, Cassia Trojahn, and Romain Hugues. 2020. Giving meaning to unsupervised EO change

detection rasters: a semantic-driven approach. In *9th ACM SIGSPATIAL International Workshop on Analytics for Big Geospatial Data (BIGSPATIAL'20)*, November 3–6, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3423336.3429347>

1 INTRODUCTION

Earth Observation (EO) is a domain that has greatly evolved in the last years thanks to large-scale Earth monitoring programs, such as the US Landsat Program¹ and the EU Copernicus Program². In particular, with the Copernicus program launched by the European Space Agency (ESA), data is collected by Earth observation satellites and combined with observation data from sensor networks on the Earth's surface. Nowadays, two types of satellites are in production, Sentinel-1 and Sentinel-2, with several other types being expected by 2030. Since 2015, Sentinel-1 and Sentinel-2 are delivering high-quality Earth images (estimated between 8TB to 10TB of data daily), providing users with free, reliable and up-to-date Earth image data and metadata. The availability of these data sources has opened the opportunity to better support existing domain-oriented applications and to boost emerging new ones, from agriculture to forestry, environmental monitoring to urban planning, climate studies and disaster monitoring. In all these domains, change detection using series of satellite images is of paramount importance [4, 23], supporting decision making, monitoring or risk management. EO change detection is the process of comparing two or more satellite images of the same area on the Earth surface taken at different points in time. The goal of this task is to identify differences between the images in areas as detailed as the image pixel. As an example, let's consider Figures 1(a) and 1(b), which depict snapshots of the same area, located on North of San Francisco Bay area in California on 2019, October 22nd and 2019, November 1st. This area has been affected by an important fire during that period, which can be hardly detected by human eye when comparing the two images. This is why imagery, and particularly change detection on images, can be a means to check precise fire impacts after they occurred, or to follow them while they are still active, even in hard-to-reach places.

Over the past few years, change detection on images has been fostered thanks to increasingly efficient deep learning algorithms,

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
BIGSPATIAL'20, November 3–6, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-8162-8/20/11...\$15.00
<https://doi.org/10.1145/3423336.3429347>

¹<https://www.usgs.gov/land-resources/nli/landsat>

²<http://www.copernicus.eu/en>

Figure 1: Sentinel-2 images of California used as input data ((a) 22/10/2019 - (b) 01/11/2019). The red rectangle indicates an area affected by the fire after the first image.

with or without supervision, dealing with different image resolutions [2, 13] and applying different levels of optimization [2, 16]. These algorithms are based on a core process that is able to identify changes in pixels between two images or a series of them. In general, they produce a raster change map as output, where each pixel (numerical) value evaluates the degree of change. However,

their output lacks context. Without additional information or context, it is almost impossible to know the nature of the events that caused the changes. In the case of the above example, the change detection algorithm outputs a raster file where the highlighted areas on the second image are evaluated as highly changed between the captures of image 1 and image 2. But without any additional information, it is hard to provide an explanation for this change. A human may know that the actual reason for this change is a fire thanks to contextual information (personal contact, pieces of news, tweets or any social network data, or fire maps from other sensor measures). Providing contextual knowledge for explaining a detected change helps domain experts analyzing those images, understanding the phenomenon, and taking the appropriate decision. Back to the example, relevant information about the changes can be the administrative units, so that one knows in which place (city and/or county) the fire took place, land cover to know the kind of affected vegetation or inhabited area, tweets or pieces of news that could confirm that the event causing the change is a fire, and temperature sensor measures, hoping that higher temperature confirms the fire activity.

This paper proposes a semantic-driven approach for characterizing Earth Observation (EO) changes detected by unsupervised deep learning algorithms. This approach relies on a semantic representation of image tiles and open data sources that provide context to the images. Its output is a knowledge graph that represents spatial and temporal features for areas affected by a high change, as well as various kinds of contextual data useful for explaining the detected changes. This output can hence be exploited for a) explaining change detection and b) validating the results from unsupervised deep learning algorithms. We applied our approach in a use case of fire monitoring. Three sources of data have been exploited (tweets, administrative units and thermal maps captured by embedded sensors) and made available through an RDF endpoint so that it could be queried to help user in better understanding the changes.

The rest of this paper is organised as follows. Section 2 discusses the main related works. Section 3 overviews our approach while Section 4 details the knowledge graph generation process. Section 5 describes the evaluation of the approach in terms of changes explanation and detection of false negatives. Finally, Section 6 concludes the paper and presents future work.

2 RELATED WORK

Image change detection. Image change detection is the study of detecting changes between two different images with the same footprint, taken at different times, with different levels of detection. Qualifying a change only based on image processing is not enough and that is why it is relevant to combine these results with data from other sources. Annotating the images with external data or with data calculated from the images themselves has been addressed in different proposals. In [21], vegetation indexes are computed from satellite image processing and exposed as RDF triples using GeoSPARQL [19]. The changes in these indexes are used for supporting forest monitoring. Close to that approach, in [18] statistical data is collected from open sources to monitor the deforestation of Amazonian rainforest, with temporal series, and translated into RDF. In [20], background information (OpenStreetMap, GeoNames,

etc.) is used to identify hot spots concerned by a wildfire, in order to refine the process and to reject hot spot candidates. The work made by [14] uses heuristics to detect changes between two satellite images, extracting extract contextual change information such as damage caused by natural hazards.

With a learning perspective, the work from [1] applies a Convolution Neural Network (CNN) to detect different classes of land cover in the satellite image (vegetation, ground, road, building, water, railroad, parking), which are enriched with semantic labels from OpenStreetMap and government sources. The study by [17] involves inserting semantic meaning into detected change areas with hypermaps, as a way of indexing the semantic information to pixels. They apply their CNN method to the TSUNAMI Panoramic Change Detection dataset and re-annotate the changed areas with semantic classes (car, building and rubble). In the same line, [10] proposes a CNN-based learning approach that simultaneously performs change detection and land cover mapping, while using the predicted land cover information to help to predict changes. Also close to our study, [22] enriches change detection over Sentinel-1-A images with event detected in media content (news, posts). Social sensing applies event detection techniques to cluster together news items and social media posts that pertain to the same real-world event and are located in the area where changes were detected.

Geo-spatial and temporal data linking. Interlinking data on EO means discovering spatial and temporal links among RDF graphs [8]. It is close to what we have done in terms of data integration. Thanks to spatial links, data from observations can be associated to tiles and then to EO images. Thanks to temporal links, temporal observations can be linked to images too. In case entities of the same nature are collected from various sources, an entity resolution algorithm can identify mappings between similar or identical spatial entities. We are concerned only by temporal and spatial relationships. The OGC introduced the notion of *geolinked data* to refer to geographically related data. In early works, geometry was not directly stored within the attribute data, but in a separate geo-spatial data-set. This option adds constraints when comparing the geometry of each entity. However, current repositories store together an RDF representations of the geometry with the RDF spatial entities. In [5], various types of geometries (point, line or polygon) have been identified, together with various tools to build an RDF representation of the geometry (like Geometry2RDF³ or TripleGeo⁴). The process from [9] compares data geometries, so that spatial data could be retrieved and interlinked on a high level of granularity. We have adopted a similar strategy in our approach, and rely on a precise comparison of the spatial component of each entity to integrate data. More recently, [1] proposes a framework in which satellite images are classified and enriched with additional semantic data in order to enable queries about what can be found at a particular location.

In order to compute links between Linked Open Data (LOD) resources with temporal, and thus event-like properties, the work in [12] uses the intervals of Allen's algebra. Their approach reduces the number of Allen's temporal relationships from 13 to 8, and optimally implements them to more quickly perform the temporal

property comparisons needed to compute temporal relationships. Another way of linking is referred to entity matching (or *entity resolution*). It is a matter of associating equivalent entities. More generally, link discovery (*entity linking*) aims to find semantic links between entities from different knowledge bases [7, 25]. According to [25], most approaches focus on the search for equivalence between entities (same labels, same names or same types), leaving other types of relationships, e.g. spatial or temporal relations, unexploited. They propose to use spatio-temporal links to improve the process. However, the spatial representation of most geo-located data is complex, taken the form of a polygon. Calculating relation between polygons in very large datasets is particularly complex and time consuming. A pre-processing step is necessary to transform the data (from RDF vocabularies, CRS, serializations, etc.) into a single model. Then, a technique of *blocking* aims to reduce the computational complexity. It consists of dividing the earth's surface into "blocks" (curved rectangles) and then evaluating the topological relations between entities based on this division. Similarly, [24] propose to discover topological links even more efficiently by indexing entities using tiles cutting the land surface into rectangles. This method speeds up the computation of topological relationships between two feature geometries in the map.

As [5], we propose an ontology that extends standard vocabularies to better represent image metadata and their contextual data as entities associated with classes having spatial (geometry) and temporal (dating) properties. To integrate these data, we first rely on their spatial dimension and like [25]. We then use the notion of ROI which corresponds to the minimum bounding box of every polygons extracted from the machine learning change algorithm output to reduce the costs of computing spatial relationships between entities representing data and images, in the sense of [24]. However, we chose to limit ourselves to the spatial relationships defined by GeoSPARQL in order to use this language to query the data. In a second step, the integration takes into account the temporal and spatial properties of the data, as detailed in the following.

3 OVERVIEW OF THE SEMANTIC-DRIVEN DATA INTEGRATION APPROACH

This section presents our approach to detect high change areas from the result of a machine learning algorithm on satellite images and to associate contextual data to them. This process produces an RDF knowledge graph (KG, in short), as shown in Figure 5 from data selected from various sources with heterogeneous formats and structures. We illustrate the approach with a running use case on fire monitoring.

3.1 Data sources

The data integration process takes as input two kinds of data organized in two different types of datasets:

- Change rasters computed from Sentinel-2 images and the corresponding tile definitions;
- Contextual data with spatial and temporal features available as open data or extracted from text, social networks, etc.

3.1.1 Change rasters and tiles.

³<https://github.com/boricles/geometry2rdf>

⁴<https://github.com/GeoKnow/TripleGeo>

Sentinel-2 tile: In order to avoid duplicating static data that would tag all the images of the same area over time, the notion of tile defined by ESA is convenient. When using a tile grid, the whole Earth surface is associated a grid where a tile represents a fixed area on this surface. Each Sentinel-2 image is associated a single tile with a surface of 100km^2 , which means that the image footprint is a tile. Tiles are linked to EO images as one of their metadata.

Change raster: A change raster is produced from a couple of Sentinel-2 images having the same tile, and captured at different dates. The tile and the dates are part of the raster metadata. It is a grid in which each pixel is associated a change value according to the change level computed by a given algorithm. Figure 2 is an example of change raster.

3.1.2 Contextual data.

Firecast dataset: Data collected from the Firecast⁵ website comes from several satellites like MODIS and VIIRS. Thermal sensors on these satellites capture Earth temperature and make it possible to identify wild fires with a confidence value. So, this dataset contains all the detected fires with their confidence index in shape file format. The fire points are updated daily which provides an NRT (Near Real Time) monitoring. Each fire point is located with coordinates and has a timestamp.

Yago2geo dataset: The second contextual data source we used here is the Yago2geo⁶ dataset, accessible via a SPARQL endpoint. This dataset gathers administrative units with their spatial data from different sources. Administrative units are geolocated with a polygon using the WKT format. We use them to geolocate related entities (such as events).

Social data: The third contextual knowledge source is the Twitter social network. The keywords in tweets help identifying the kind of change that has occurred in the tile footprint during the period of the change raster. The tweets are extracted from the tweet collection captured in our lab, IRIT, on the OSIRIM⁷ platform. This platform is storing a corpus containing 1 percent of the global tweets since 2016. It collects between 20 and 30 tweets per seconds without any restriction criteria. In the current keyword extraction process, place name ambiguity is not treated. In order to deal with English stop words we used the open source module NLTK (Natural Language Toolkit) provided by Python.

3.2 Main steps of the process

The overall KG generation is illustrated in Figure 3. It takes a change raster and the various contextual datasets as input, and generates a set of RDF triples. We give here an overview of the three main steps of the process, and we will set them out in the next section:

Step 1: Processing image change raster Given a change raster file, we get the image footprint (the tile geometry) and the raster period. Then, we identify on the raster areas where

Figure 2: Interpretation of a change raster with the following change levels: No Change (blue), Low Change (green), Middle Change (yellow), and High Change (red).

changes are the most significant. We call these areas *Regions of Interest* (ROIs).

Step 2: Retrieving contextual data Considering the tile geometry and the period of concern of the images, contextual data is retrieved from each of the data sources introduced above.

Step 3: Generating the knowledge graph A set of transformation rules are applied to extract the appropriate data from the different data sources and to create an RDF graph. Several RDF graphs are conceptually linked together into one, as soon as a node identifier (URI) appears in each of them. Thus the set of all the RDF sub-graphs generated from a raster forms a KG (Figure 5). It can be used to retrieve the number of ROIs and their size, the number of firepoints and the name of the cities located in these ROIs, and keywords extracted from tweets dealing with these cities (Figure 5.2).

3.3 Semantic vocabulary

We have defined a vocabulary to represent contextual data sources and rasters. It reuses standard vocabularies, namely GeoSPARQL (for geospatial data) and OWL-Time [15] (for temporal series), with respectively the *geo* and *time* prefixes.

GeoSPARQL [19] is an OGC standard which allows to represent an entity of the real world (called *geo:Feature*) and geometries (called *geo:Geometry*). The first one represents spatial entities while the second represents all geometric forms defined on a spatial coordinate reference system. An entity is associated to its geometries by the *geo:hasGeometry* property. GeoSPARQL provides topological relations and functions to link spatial objects (intersects, touches, etc.). Geo-spatial entities, i.e. instances of *geo:Feature*, described with a specific vocabulary (a specific prefix), are tiles (*o-grid*), ROIs (*o-change*) and administrative units (*o-admin*).

OWL Time [15] is recommended by the W3C for modeling temporal entities (instants and intervals) and expressing topological

⁵<https://firecast.conservation.org/>

⁶<http://yago2geo.di.uoa.gr/>

⁷<https://osirim.irit.fr/site/>

Figure 3: Pipeline for generating a knowledge graph from a change raster.

relations as defined in the theory of Allen between them (before, after, etc.). It is used to describe temporal series. Rasters (*o-change*), social data (*o-socialData*), and fire points (*o-firecast*) are instances of *time:TemporalEntity*. The temporal dimension of a ROI is provided by its raster (*o-change:hasHighChangePolygon* property). While the spatial dimension of a raster and a firepoint is provided by the tile to which they are linked (property *o-change:hasChange* for rasters and property *o-firecast:hasFirePoint* for fire points), the spatial dimension of social data is the geometry of the administrative unit associated with it using a property *o-socialData:hasSocialData*.

4 FROM CHANGE RASTERS TO RDF GRAPHS STEP BY STEP

In the following, we will detail the main steps of the workflow introduced above.

4.1 ROI processing

This step consists of transforming a change raster to a shapefile containing all the polygons corresponding to the ROIs. Similar pixels are first grouped. We used the open source library GDAL and the polygonize algorithm in order to produce a shapefile containing polygons formed by highest changes values from the raster. As the change raster contains values between 0 for the lowest changes and 1 for the most important changes, we decided to polygonize the pixels with a value above 0.66. Having these polygons represented in a shapefile, we then represented each polygon with its envelope. When having several envelopes intersecting themselves, they are merged into one polygon. Only pixels having a high change value (above a predefined threshold) are considered in order to process

significant changes. This process gives many polygons, several of them having only one pixel. Several iterations are thus made in order to group close polygons: when two polygons touch each other they become one and the new shape is the bounding box grouping these two polygons. At the end all polygons are disjoint and correspond to ROIs (Pink areas in Figure 4). Using the size of these polygons we are able to identify where are the most important changes on the tile.

Figure 4: Representation of high change polygons (ROI) on a Sentinel 2 tile

This is an example of the corresponding RDF graph produced. Entities can be identified using their URI:

- ML_Change_: a change raster
- ML_EOData_: raster metadata
- Tile_: a Sentinel 2 tile
- highChange_: a ROI

```
# the change raster
g-ch:ML_EOData_T10SEH_20191022T214432_20191101T213037
  o-change:hasSource
    "S2B_MSIL2A_20191022T185429_T10SEH..."^^xsd:string .

g-ch:ML_EOData_T10SEH_20191022T214432_20191101T213037
  o-change:hasSource
    "S2B_MSIL2A_20191101T185529_T10SEH..."^^xsd:string .

g-ch:ML_EOData_T10SEH_20191022T214432_20191101T213037
  time:hasBeginning g-ch:Instant_20191022T214432;
  time:inXSDDateTimeStamp
    "2019-10-22T21:44:32Z"^^xsd:dateTimeStamp .

g-ch:ML_EOData_T10SEH_20191022T214432_20191101T213037
  time:hasEnd g-ch:Instant_20191101T213037 ;
  time:inXSDDateTimeStamp
    "2019-11-01T21:30:37Z"^^xsd:dateTimeStamp .

g-grid:Tile_T10SEH o-change:hasChange
  g-ch:ML_Change_T10SEH_20191022T214432_20191101T213037 .

g-ch:ML_Change_T10SEH_20191022T214432_20191101T213037
  o-change:calculatedFrom
    g-ch:ML_EOData_T10SEH_20191022T214432_20191101T213037 .

# a ROI
g-ch:ML_Change_T10SEH_20191022T214432_20191101T213037
  o-change:hasHighChangePolygon
    g-ch:highChange_T10SEH_20191022_20191101_roi0 .

g-ch:highChange_T10SEH_20191022_20191101_roi0
  geo:hasGeometry
    g-ch:highChange_T10SEH_20191022_20191101_roi0_geo .
g-ch:highChange_T10SEH_20191022_20191101_roi0_geo geo:asWKT
  "POLYGON (...)"^^geo:wktLiteral .
```

4.2 Event selection

From the firecast website, we have downloaded the data corresponding to the Earth area and the period covered by the raster. The period starts from the date of the first image and the end of the period is the date of the second image taken for the change detection. We do not take into account other dates between this period as we took the Sentinel 2 images with the closest date from each other. In order to avoid false positives while keeping a significant number of firepoints, we have only retrieved those having a confidence index greater than 0.65. The example in the following corresponds to the RDF graph produced for one fire point:

```
g-firecast:FirePoint_108566952 geo:hasGeometry
  g-firecast:FirePoint_108566952_geo .
g-firecast:FirePoint_108566952_geo geo:asWKT
  "POINT (...)"^^geo:wktLiteral .

g-firecast:FirePoint_108566952 time:hasTime
```

```
g-firecast:instant_1571644800 .
g-firecast:instant_1571644800 time:inXSDDateTimeStamp
  "2019-10-21T10:00:00Z"^^xsd:dateTimeStamp .

g-firecast:FirePoint_108566952
  o-firecast:hasConfidence "75.0"^^xsd:decimal .
g-firecast:FirePoint_108566952
  o-firecast:hasType "VIIRS"^^xsd:string .

g-grid:Tile_T10SEH o-firecast:hasFirePoint
  g-firecast:FirePoint_108566952 .
```

4.3 Administrative unit selection

In order to have the name of the different cities located in the raster tile, we retrieve in Yago2geo those with a geometry that intersects the tile geometry using a SPARQL query. This is relevant for the next step consisting in retrieving data from the social networks. Yago2geo contains an important part of all the data of Open Street Map⁸ and Global Administrative Areas (GADM)⁹.

The RDF graph produced for each administrative unit looks as follows:

```
g-grid:Tile_T10SEH o-admin:hasAdminUnit
  g-admin:Richmond-California .
g-admin:Richmond-California owl:sameAs
  <http://yago-knowledge.org/resource/Richmond,_California> .
```

4.4 Keyword selection

From the social data source, we retrieved all the keywords of available tweets related to the raster tile footprint and during the period of the analysed images. The goal here is to fetch keywords from the topic of tweets. So we decided to search all the tweets that contain the name of the administrative areas identified at step 3 and between the two defined dates. The reason why we use the names of the administrative units rather than the tweet locations is that some people can tweet about a particular area being located in another area. We chose to keep the 5 keywords with the highest frequency for an administrative area. To do so, we used the NLTK (Natural Language Toolkit) python module.

In the running example, two keywords (“evacuation” and “#kin-cadefire”) could be collected from tweets between 2019-10-23 and 2019-11-01 about the city of Calistoga. The example in the following corresponds to the generated RDF graph:

```
g-admin:Calistoga-California o-socialData:hasSocialData
  g-sd:SocialData_Calistoga-California_20191023_20191101 .

g-sd:Instant_20191023 time:inXSDDate
  "2019-10-23"^^xsd:date .
g-sd:SocialData_Calistoga-California_20191023_20191101
  time:hasBeginning g-sd:Instant_20191023 .

g-sd:Instant_20191101 time:inXSDDate
  "2019-11-01"^^xsd:date .
g-sd:SocialData_Calistoga-California_20191023_20191101
  time:hasEnd g-sd:Instant_20191101 .

g-sd:Keyword_evacuation o-socialData:hasValue
```

⁸<https://www.openstreetmap.org/>

⁹<https://gadm.org/>

```

    "evacuation"^^xsd:string .
g-sd:Keyword_evacuation o-socialData:hasFrequency
    "234"^^xsd:decimal .
g-sd:SocialData_Calistoga-California_20191023_20191101
    o-socialData:hasKeyword g-sd:Keyword_evacuation .

g-sd:Keyword_kincadefire o-socialData:hasValue
    "#kincadefire"^^xsd:string .
g-sd:Keyword_kincadefire o-socialData:hasFrequency
    "229"^^xsd:decimal .
g-sd:SocialData_Calistoga-California_20191023_20191101
    o-socialData:hasKeyword g-sd:Keyword_kincadefire .

```

4.5 Generating a knowledge graph

The process that generates a knowledge graph is similar for all data sources. Given a source dataset, once data has been extracted according to spatial and temporal features, it is converted to JSON format. Then, a Python script uses a template to generate RDF triples from the JSON file. A *template* defines the mapping between a JSON schema and an RDF model to be instantiated. In our case, JSON schemes match the data structure of the data sources listed in Section 3.1.2; and the model is the vocabulary presented in Section 3.3. We wrote one specific template in Turtle format for each data source. We reused and evolved the mapping template and processing mechanism of previous work [3]. These templates and functions help to perform sophisticated mapping operations that are not possible in alternative approaches like RML rules. Moreover, they make it easy to enrich the knowledge graph with new data, or to generate a new graph if either the data source schema or the modelling vocabulary are changed. One just has to adapt the template and to run the process with the new template and data. As an illustration, the following template applies to the fire points extracted from the Firecast source.

```

dummy      a                o-firecast:FirePoint .
dummy      geo:hasGeometry  dummy_geo .
dummy_geo  geo:asWKT        valueToWktLiteral($.geometry) .
dummy      time:hasTime     stringValueToTimeInstant($.date) .
dummy      o-firecast:hasConfidence
    valueToDecimalLiteral($.confidence) .
dummy      o-firecast:hasType stringToLiteral($.type) .

```

5 EVALUATION

Supervised machine learning approaches usually require the user to annotate a large set of training examples, i.e. changes in the perspective of change detection, like in [11]. As shown in Figures 1(a) and 1(b), some cases of change may be hard to be detected by humans. In an unsupervised setting, automatic semantic annotations are of paramount importance for helping users to evaluate learning results. We argue here that semantic annotations help a) for explaining change detection and b) for validating results from unsupervised deep learning algorithms. We validate these hypotheses through an empirical experiment with changes detected on Sentinel-2 images (Section 5.1) and the data sources listed in Section 3.1.2 to provide semantic annotations. We performed the process detailed in Section 4 to generate RDF triples which can be queried either to explain the changes to end-users (Section 5.2)

or to validate the changes detected by the unsupervised learning process (Section 5.3).

5.1 Change detection dataset

We have experimented our process on change raster files produced by Thales Alenia Space from Sentinel-2 images thanks to an unsupervised machine learning algorithm [6]. This algorithm relies on TensorFlow and requires no previous annotation. The model has been trained on collections of pairs of Sentinel-2 images with the same tile captured at different dates, which makes it generic enough to process any Sentinel-2 pair of images of the same area. For each input pair of images, its output is a one band GeoTiff raster file where each pixel is associated a change value according to the change level computed by the algorithm. The change value varies from 0 (no change detected) to 1 (high change detected). We only consider pixels having a value between 0.66 and 1.

5.2 Explaining change detection

The following SPARQL query retrieves all the ROIs of tiles, the number of linked fire points and the keywords of the cities in the ROI for the period. The first step of the query collects the change raster data. This raster gives us the raster date and location as well as its Sentinel-2 tile. The next step of the query is dedicated to fetch all the ROIs on this change raster and their geometry. For each ROI it collects all the fire points that intersects the geometry and the time period. We chose to count the number of these fire points. The other contextual data we can have on a high change polygon is the administrative units that intersect the geometry. Labels of administrative units are important because they are used to get keywords from social data (Twitter data source) about these administrative units. Retrieving the keywords associated to the administrative unit and their frequency is the last step of the query.

```

SELECT ?tile ?start ?end ?roi
(count(distinct ?firepoint) as ?nbFirePointROI) ?adminUnitName
?keywordAdminUnit ?keywordFrequency

```

WHERE{

```

#Step 1 - Change raster data
?MLchange o-change:calculatedFrom ?ML_EO_Data .
?ML_EO_Data o-change:hasSource ?image .
?ML_EO_Data time:hasBeginning/time:inXSDDateTimeStamp ?start .
?ML_EO_Data time:hasEnd/time:inXSDDateTimeStamp ?end .
?MLchange ^o-change:hasChange ?tile .

```

```

#Step 2 - ROI list
?MLchange o-change:hasHighChangePolygon ?roi .
?roi geo:hasGeometry/geo:asWKT ?roiWkt .

```

```

#Step 3 - Firepoints
?tile o-firecast:hasFirePoint ?firepoint .
?firepoint geo:hasGeometry/geo:asWKT ?firepointWKT .
?firepoint time:hasTime/time:inXSDDateTimeStamp ?firepointTime .
FILTER(bif:st_intersects (?firepointWKT, ?roiWkt)) .
FILTER(?firepointTime > ?start && ?firepointTime < ?end)

```

```

#Step 4 - Administrative units
?tile o-admin:hasAdminUnit ?adminUnit .
?adminUnit rdfs:label ?adminUnitName .
?adminUnit geo:hasGeometry/geo:asWKT ?adminWKT .
FILTER(bif:st_intersects (?adminWKT, ?roiWkt)) .

```

```

#Step 5 - Keywords from Twitter
?adminUnit o-socialData:hasSocialData ?socialData .

```

Figure 5: Knowledge graph (KG) model of a change raster.

```

?adminSocialData time:hasBeginning/time:inXSDDateTimeStamp
?socialDataStart .
?adminSocialData time:hasEnd/time:inXSDDateTimeStamp
?socialDataEnd .
?socialData o-socialData:hasKeyword ?socialDataKeyword .
?socialDataKeyword o-socialData:hasValue ?keywordAdminUnit .
?socialDataKeyword o-socialData:hasFrequency ?keywordFrequency .
FILTER(?socialDataStart > ?start && ?socialDataEnd < ?end)
}
ORDER BY ?adminUnitName

```

From the result of this query on the RDF triples built from a change raster, we can get the ROI that contains the highest number of fire points in the area of this raster and the keywords associated to the administrative unit that intersects with this ROI. Keywords can contribute to determine if there is a fire in the ROI. In Figure 6, we can see that the ROI number 273 contains 2167 fire points during the raster time period (from 2019-10-22 to 2019-11-01). The word “fire” is one of the keywords detected at this period for the city of Healdsburg in California. Hence, contextual data explains that the change observed on this area has been caused by a fire. This data documents properly the kind of event that has caused a detected high change.

5.3 Enriching results from unsupervised learning

The RDF triples can also contribute to validate the results from the machine learning algorithm. To do so, we cross-checked the results obtained from the different sources on the same tile and period of time as in the previous section. The results from deep learning

process show a high change on the image. The Firecast data show also many fire points in the area that explain that the change is due to a fire. This information is also corroborated by the data collected from tweets, that report a forest fire called Kincade Fire¹⁰ in this area of California during this period.

Contrariwise, one can observe a fire in a smaller area in the lower right part of the image, as indicated by a grouping of fire points (Figure 7). However, this fire has not been detected as a significant change from the change raster because there are not enough pixels with a value above 0.66 to detect an abnormal change. Nevertheless, our processing chain was able to detect multiple fire points thanks to the Firecast data (see Figure 4). In that way, our approach is able to identify false negative areas of the change. Moreover, it corroborated the changes detected by the unsupervised algorithm. In a further work, RDF data could be more tightly coupled with the change learning process to improve its results, for instance by providing feedback or training examples.

6 CONCLUSIONS

Machine learning algorithms for change detection in images are able to identify changes at pixel level. Pixels with similar change values can be aggregated to detect change areas. In unsupervised approaches, no information is available about the phenomenon that caused the change. An interpretation of the “meaning” of each change would require contextual information to be integrated either in the machine learning process or after the change areas are identified. We proposed a semantic-based support to the second

¹⁰<https://www.fire.ca.gov/incidents/2019/10/23/kincade-fire/>

tile	start	end	roi	nbFirePointROI	adminUnitName	keywordAdminUnit	keywordFrequency
http://melodi.irit.fr/iod/grid/Tile_T10SEH	2019-10-22T21:44:32	2019-11-01T21:30:37	http://melodi.irit.fr/iod/change/highChange_T10SEH_20191022_20191101_roi273	2167	"Healdsburg"	"fire"	6228
http://melodi.irit.fr/iod/grid/Tile_T10SEH	2019-10-22T21:44:32	2019-11-01T21:30:37	http://melodi.irit.fr/iod/change/highChange_T10SEH_20191022_20191101_roi273	2167	"Healdsburg"	"evacuation"	8024
http://melodi.irit.fr/iod/grid/Tile_T10SEH	2019-10-22T21:44:32	2019-11-01T21:30:37	http://melodi.irit.fr/iod/change/highChange_T10SEH_20191022_20191101_roi273	2167	"Healdsburg"	"windsor"	7207
http://melodi.irit.fr/iod/grid/Tile_T10SEH	2019-10-22T21:44:32	2019-11-01T21:30:37	http://melodi.irit.fr/iod/change/highChange_T10SEH_20191022_20191101_roi273	2167	"Healdsburg"	"healdsburg"	15324
http://melodi.irit.fr/iod/grid/Tile_T10SEH	2019-10-22T21:44:32	2019-11-01T21:30:37	http://melodi.irit.fr/iod/change/highChange_T10SEH_20191022_20191101_roi273	2167	"Healdsburg"	"#healdsburg"	4779
http://melodi.irit.fr/iod/grid/Tile_T10SEH	2019-10-22T21:44:32	2019-11-01T21:30:37	http://melodi.irit.fr/iod/change/highChange_T10SEH_20191022_20191101_roi273	2167	"Healdsburg"	"#kincadeville"	6892

Figure 6: Result of SPARQL query of a change raster and contextual data

Figure 7: Fire points as image semantic annotations.

option, i.e. to provide contextual information that can document changes. More precisely, we make it possible to confirm that automatically detected changes are related to a specific type of event. We addressed this problem for Earth observation images captured by Sentinel-2 satellites, and changes detected by unsupervised deep learning algorithms. We defined a generic process that takes dated and geolocated changes in raster files as input, various open dated and geolocated data sources (events of interest, territorial units, social network short messages, etc.) as contextual information, to produce a knowledge graph. An important concept in the data integration vocabulary is the notion of tile, that plays a major role to structure data according to its location on Earth. We showed that by querying this graph one can get (i) contextual information about changes, like the name of the closest city, the event that caused the change or its date, and (ii) a kind of confirmation and assessment of the confidence that a change actually occurred.

We plan to carry out new experiments to confirm the generality of the integration process, using similar or new types of events

(i.e. urbanization, migration flows, ...) and other types of contextual data. We will also investigate how to reuse the RDF data to improve the change detection process.

Acknowledgements. This work benefits from two financial supports: an H2020 grant for the CANDELA project under convention number 776193; and the CIFRE PhD. scholar agreement ANRT number 2017/1399 involving Thales Alenia Space and CNRS.

REFERENCES

- [1] Marjan Alirezaie, Andrey Kiselev, Martin Långkvist, Franziska Klügl, and Amy Loutfi. An ontology-based reasoning framework for querying satellite images for disaster monitoring. *Sensors (Switzerland)*, 17, 11 2017.
- [2] Arabi Mohammed El Amin, Qingjie Liu, and Yunhong Wang. Convolutional neural network features based change detection in satellite images. In Xudong Jiang, Guojian Chen, Genci Capi, and Chiharu Ishii, editors, *First International Workshop on Pattern Recognition*, volume 10011, pages 181 – 186. International Society for Optics and Photonics, SPIE, 2016.
- [3] Helbert Arenas, Nathalie Aussenac-Gilles, Catherine Comparot, and Cassia Trojahn. Semantic integration of geospatial data from earth observations. In *Knowledge Engineering and Knowledge Management - EKAW 2016 Satellite Events*, pages 97–100, Bologna (It), 2016. Springer.
- [4] Anju Asokan and J. Anitha. Change detection techniques for remote sensing applications: a survey. *Earth Science Informatics*, 12(2):143–160, Jun 2019.
- [5] Ghislain Auguste Atemezang. *Publishing and consuming geo-spatial and government data on the semantic web*. PhD thesis, Thesis, 04 2015.
- [6] M. Aubrun, A. Troya-Galvis, M. Albughadi, R. Hugues, and M. Spigai. Unsupervised learning of robust representations for change detection on sentinel-2 earth observation images. In *Proc. of the 13th Int. Symp. on Environmental Software Systems*, Wageningen (NL), 2020.
- [7] Sören Auer, Jens Lehmann, and Axel-Cyrille Ngonga Ngomo. *Introduction to Linked Data and Its Lifecycle on the Web*, pages 1–75. 2011.
- [8] Luis Manuel Vilches Blázquez, Victor Saquicela, and Óscar Corcho. Interlinking geospatial information in the web of data. In *Bridging the Geographic Information Sciences - International AGILE'2012 Conference*, Avignon, France, April, 24–27, 2012, pages 119–139, 2012.
- [9] Luis Manuel Vilches Blázquez, Boris Villazón-Terrazas, Óscar Corcho, and Asunción Gómez-Pérez. Integrating geographical information in the linked digital earth. *International Journal of Digital Earth*, 7(7):554–575, 2014.
- [10] Rodrigo Caye Daudt, Bertrand Le Saux, Alexandre Boulch, and Yann Gousseau. High resolution semantic change detection. *CoRR*, abs/1810.08452, 2018.
- [11] C. Dumitru, G. Schwarz, and M. Datcu. Land cover semantic annotation derived from high-resolution sar images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 9(6):2215–2232, June 2016.
- [12] Kleanthi Georgala, Mohamed Ahmed Sherif, and Axel-Cyrille Ngonga Ngomo. An efficient approach for the generation of allen relations. In *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, pages 948–956, 2016.
- [13] Maoguo Gong, Hailun Yang, and Puzhao Zhang. Feature learning and change feature classification based on deep learning for ternary change detection in sar images. *ISPRS Journal of Photogrammetry and Remote Sensing*, 129:212 – 225, 2017.

- [14] S. Hashimoto, T. Tadono, M. Onosato, M. Hori, and T. Moriyama. A framework of ontology-based knowledge information processing for change detection in remote sensing data. In *2011 IEEE International Geoscience and Remote Sensing Symposium*, pages 3927–3930, 2011.
- [15] J. R. Hobbs and F. Pan. An ontology of time for the semantic web. *ACM Transactions on Asian Language Information Processing*, 3:66–85, 2004.
- [16] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*, 2014.
- [17] Hirokatsu Kataoka, Soma Shirakabe, Yudai Miyashita, Akio Nakamura, Kenji Iwata, and Yutaka Satoh. Semantic change detection with hypermaps. *CoRR*, abs/1604.07513, 2016.
- [18] Tomi Kauppinen, Giovana Mira de Espindola, Jim Jones, Alber Sanchez, Benedikt Graeler, and Thomas Bartoschek. Linked brazilian amazon rainforest data. *Semantic Web Journal*, 5, 01 2013.
- [19] Dave Kolas, Matthew Perry, and John Herring. Getting started with GeoSPARQL. Technical report, OGC, 2013.
- [20] Kostis Kyzirakos, Manos Karpathiotakis, George Garbis, Charalampos Nikolaou, K Bereta, Ioannis Papoutsis, T Herekakis, Dimitrios Michail, Manolis Koubarakis, and Kontoes Charalabos. Wildfire monitoring using satellite images, ontologies and linked geospatial data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 24, 01 2014.
- [21] Antonio Jesús Pérez Luque, Ramón Pérez Pérez, Francisco Javier Bonet, and P.J. Magaña. An ontological system based on modis images to assess ecosystem functioning of natura 2000 habitats: A case study for quercus pyrenaica forests. *Int. Journal of Applied Earth Observation and Geoinformation*, 37:142–151, 05 2015.
- [22] Nikiforos Pittaras, George Papadakis, George Stamoulis, Giorgos Argyriou, Efi Karra Taniskidou, Emmanouil Thanos, George Giannakopoulos, Leonidas Tsekouras, and Manolis Koubarakis. Geosensor: Semantifying change and event detection over big data. In *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing, SAC '19*, pages 2259–2266, New York, NY, USA, 2019. ACM.
- [23] R. J. Radke, S. Andra, O. Al-Kofahi, and B. Roysam. Image change detection algorithms: a systematic survey. *IEEE Transactions on Image Processing*, 14(3):294–307, March 2005.
- [24] Mohamed Ahmed Sherif, Kevin Dreßler, Panayiotis Smeros, and Axel-Cyrille Ngonga Ngomo. Radon - rapid discovery of topological relations. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 175–181, 2017.
- [25] Panayiotis Smeros and Manolis Koubarakis. Discovering spatial and temporal links among RDF data. In *Proceedings of the Workshop on Linked Data on the Web, LDOW 2016, co-located with 25th International World Wide Web Conference (WWW 2016)*, 2016.