

An Ontology-based approach for making Machine Learning systems Accountable

Iker Esnaola-Gonzalez^{a,*}

^a *TEKNIKER, Basque Research and Technology Alliance (BRTA), Iñaki Goenaga 5, 20600 Eibar, Spain*
E-mail: iker.esnaola@tekniker.es

Editors: Michel Dumontier, Maastricht University, The Netherlands; Sabrina Kirrane, Vienna University of Economics and Business, Austria; Oshani Seneviratne, Rensselaer Polytechnic Institute, USA

Solicited review:

Open review:

Abstract. Although the maturity of the Artificial Intelligence technologies is rather advanced nowadays, its adoption, deployment and application is not as wide as it could be expected, mainly due to the lack of trust of users in the Artificial Intelligence systems. The explainable Artificial Intelligence (XAI) has emerged as a way of addressing this lack of trust. However, the explainability of the systems is necessary but far from sufficient for such a goal. Accountability, is another relevant factor to advance in this regard, as it enables discovering the causes that derived a given decision or suggestion made by an Artificial Intelligence system. In this article, the use of ontologies is conceived as the way for making Machine Learning systems accountable, as they offer conceptual modelling capabilities to describe a domain of interest, as well as formality and reasoning capabilities. The feasibility of the proposed approach has been demonstrated in a real-world scenario and it is expected to pave the way towards unlocking the full potential of Semantic Technologies for achieving trustworthy AI systems.

Keywords: Accountability, Ontology, Artificial Intelligence, Machine Learning

1. Introduction

Even though the maturity of the Artificial Intelligence (AI) technologies is rather advanced nowadays, according to McKinsey¹, its adoption, deployment and application is not as wide as it could be expected. This could be attributed to many barriers including cultural ones [1], but above all, the lack of trust of potential users in such AI systems. According to [2], AI trustworthiness can be defined as “the extent to which a user is confident in, and willing to act on the basis of, the recommendations, actions, and decisions of an artificially intelligent decision aid”.

[3, 4] studied the different factors that affect the users’ trustworthiness on AI systems. Some of these

factors comprise the so called Explainable Artificial Intelligence (XAI), which according to [5] refers to the “techniques that enable human users to understand, appropriately trust, and effectively manage the emerging generation of artificially intelligent partners”. XAI was intensively studied from the 1970s to the 1990s [6], although a resurgence of the topic has been seen recently due to the current technological advancements in the various fields of the AI [7].

According to [8], during the past decade, there has been an increase on AI systems based on black-box models, that is, models that hide their internal logic to the user. Since this lack of explanation constitutes both a practical and an ethical issue, the explainability of AI systems has been targeted by using transparent or white-box models, and by employing post-hoc explainability techniques for enhancing the interpretability of black-box models [9]. However, the explainability of AI systems is necessary but far from

* Corresponding author. E-mail: iker.esnaola@tekniker.es.

¹ <https://www.mckinsey.com/featured-insights/artificial-intelligence/ai-adoption-advances-but-foundational-barriers-remain>

1 sufficient for understanding them and holding them ac- 1
2 countable [10, 11]. Therefore, in order to develop trust- 2
3 worthy AI systems, not only should they be explain- 3
4 able, but also accountable. As a matter of fact, accord- 4
5 ing to [12], the ability to explain the inner workings, 5
6 the results and the causes of failures of AI systems to 6
7 users, regulators and citizens is critical to achieve trust 7
8 and accountability. 8

9 The accountability can be defined as the ability 9
10 to determine whether a decision was made in accord- 10
11 dance with procedural and substantive standards and 11
12 to hold someone responsible if those standards are not 12
13 met [11]. This means that with an accountable AI sys- 13
14 tem, the causes that derived a given decision can be 14
15 discovered, even if its underlying model's details are 15
16 not fully known or must be kept secret. 16

17 Therefore, it seems reasonable to consider that the 17
18 adequate representation of data, processes and work- 18
19 flows involved in AI systems could contribute to make 19
20 them accountable. There are a variety of technologies 20
21 that offer conceptual modelling capabilities to describe 21
22 a domain of interest, but only ontologies combine this 22
23 feature with Web compliance, formality and reasoning 23
24 capabilities [13]. 24

25 In this article, an ontology-based approach is pro- 25
26 posed towards achieving the accountability of Machine 26
27 Learning systems. The rest of the article is structured 27
28 as follows. Section 2 presents the related work. In Sec- 28
29 tion 3 the role of ontologies is discussed. The proposed 29
30 ontology-based approach is presented in Section 4 and 30
31 demonstrated in a real-world use case in Section 5. Fi- 31
32 nally, conclusions of this work are shown in Section 6. 32
33 33

34 2. Related Work 34

35 The usage of Semantic Technologies towards the 35
36 achievement of Trustworthy AI has not been heavily 36
37 researched in the literature, so their full potential is not 37
38 exploited yet. 38

39 [14] provides a literature-based overview of the 39
40 usage of Semantic Technologies alongside Machine 40
41 Learning methods in order to facilitate their explain- 41
42 ability. According to the reviewed literature, the main 42
43 role of the Semantic Technologies is, on the one hand, 43
44 to make Neural Networks explainable, and on the other, 44
45 to create explainable embeddings with knowl- 45
46 edge graphs. As for the domains of application, the 46
47 healthcare domain has attracted a lot of attention, al- 47
48 though they are also present in the entertainment or 48
49 commercial field. 49
50 50
51 51

1 [15] presents an approach for creating more under- 1
2 standable post-hoc explanations of decision tree algo- 2
3 rithms. In this approach, ontologies that model the con- 3
4 cerned domain knowledge are used in the process of 4
5 generating such explanations. Results showed that de- 5
6 cision trees generated with the support of domain on- 6
7 tologies are more understandable than those generated 7
8 without them. The downside of this approach is that 8
9 the used ontologies are manually created ad-hoc for 9
10 each problem, which definitely hinders its usability. 10

11 [16] proposes an explanation ontology that can be 11
12 used by designers to support the generation of different 12
13 explanation types into their AI-enabled systems. Nine 13
14 different explanation types are identified, each with 14
15 different needs, and the proposed ontology can encode 15
16 them as OWL restrictions. This provides a means for 16
17 system designers to translate their user requirements 17
18 gathered from user studies to explanations that can be 18
19 generated by their systems. 19

20 Doctor XAI is presented in [17], an ontology- 20
21 based approach for producing post-hoc explanations of 21
22 black-box sequential data classification methods. The 22
23 application of the approach is focused on the medical 23
24 domain, but since the method is agnostic with regards 24
25 to the black box model, the possible applications cover 25
26 several scenarios where a sequence of events linked 26
27 to ontology concepts can be identified, including an 27
28 online market basket analysis or Wikipedia user be- 28
29 haviour forecast. 29

30 The existing approaches, limitations and opportuni- 30
31 ties for knowledge graphs in XAI are analysed in [18]. 31
32 Furthermore, in this article knowledge graphs are en- 32
33 visioned to bring XAI to the right level of semantics 33
34 and interpretability in different AI fields, ranging from 34
35 computer vision to natural language processing. 35
36 36

37 To the extent of knowledge of author, so far, the 37
38 main focus of the usage of Semantic Technologies has 38
39 been placed on explainability, although accountability 39
40 is considered a key requirements that should be met to 40
41 achieve trustworthy AI systems [19, 20]. [21] makes 41
42 a first contribution on the usage of ontologies to sup- 42
43 port the accountability of Machine Learning systems, 43
44 proposing a method to know which predictive model 44
45 was responsible for making a given forecast, but also, 45
46 to understand where such forecasts come from, that is, 46
47 which is their underlying rationale. However, many as- 47
48 pects that could contribute to making the systems ac- 48
49 countable remain unaddressed, such as the description 49
50 of the procedure followed to develop the predictive 50
51 models. 51

All this evidence reinforces the discourse that the Semantic Technologies could play a more important role in the achieving trustworthy AI systems in general, and in solving the accountability challenge in particular.

3. On the Role of Ontologies

The term ontology comes from the Greek *ontos* (being) and *logos* (word), and it was first used in philosophy in the nineteenth century to define the study of the nature of being, existence, or reality, as well as the basic categories of being and their relations [22]. In computer and information science field, an ontology can be understood as "a formal, explicit specification of a shared conceptualisation" as defined by [23]. Therefore, ontologies appear as a way to describe and represent a certain phenomenon, topic, or subject area through the description of classes, properties and instances (also known as individuals).

According to [24] ontologies can be viewed as a spectrum of detail in their specification. Catalogues, which consist of a finite list of terms used for expressing knowledge of information, are placed in the lowest end of the spectrum. This list of terms may not have descriptions at all, and their meaning can only be estimated because they are chosen from natural language. Likewise, there are no formal relations expressed between these terms. Therefore, more often than not, such specifications are not referred to as ontologies. When at least one formal relation is defined and used between terms, the concept "ontology" can be used to refer to such a catalogue. From this point onward, there are languages that provide sets of constructs to describe an ontology, such as frames or simplified logic. As the specificity increases, the precision and the ability to use tools to automatically integrate systems also increases. However, the cost of building and maintaining a metadata registry increases accordingly.

Furthermore, [25] categorised different ontology types according to their level of generality:

- Top-level ontologies (often referred to as upper ontology or foundation ontology, general, or cross domain ontology) represent very general concepts which are independent of a specific domain or problem such as time, space and events.
- Domain ontologies describe fundamental concepts according to a generic domain and specialise terms introduced in top-level ontology.

- Task ontologies describe fundamental concepts according to a general activity or task and specialise the terms introduced in top-level ontologies.
- Application ontologies are specialised ontologies focused on a specific task and domain. They are often a specialisation of both task and domain ontologies, and they also often specify roles played by domain entities for specific activity.

3.1. Ontology development

The final quality of an ontology is directly influenced by its design and implementation, therefore, the use of well-founded ontology development methodologies such as NeOn presented by [26] or the Linked Open Terms (LOT) proposed by [27] is recommended. Furthermore, as stated by [28], the use and combination of Ontology Design Patterns (ODP) is conceived as a suitable option when developing ontologies, due to the great flexibility provided which allows a proper segmentation of the intended conceptualisation. An ODP aims at addressing recurrent ontology design problems which may arise in ontology development processes [29]. According to [30], ODPs should be extendable but self-contained, minimise ontological commitments to foster reuse, address one or more explicit requirements (such as use cases or competency questions), be associatable to an ontology unit test, be the representation of a core notion in a domain of expertise, be alignable to other patterns, span more than one application area or domain, address a single invariant instead of targeting multiple recurring issues at the same time, follow established modelling best practices, and so forth.

3.2. Reusing existing ontologies

According to [31], the reuse of ontological resources built by others and that have already reached some degree of consensus, is good practice in ontology development processes. Additionally, the W3C's Data on the Web Best practices² states that reusing an existing ontology not only captures and facilitates consensus in communities, but it also increases interoperability and reduces redundancies. Furthermore, this practice brings other important benefits:

²<https://www.w3.org/TR/dwbp/>

- 1 – It increases the quality of the applications reusing
2 ontologies, as they become interoperable and they
3 are provided with a deeper, machine-processable
4 and commonly agreed-upon understanding of the
5 underlying domain of interest.
- 6 – It reduces the costs related to ontology develop-
7 ment because it avoids the reimplementations of
8 ontological components which are already avail-
9 able on the Web and can be directly (or after some
10 additional customisation tasks) integrated into the
11 target ontology.
- 12 – It may improve the quality of the reused ontolo-
13 gies, as these are continuously revised and evalu-
14 ated by various parties through reuse.

15 As any other task in the ontology development pro-
16 cess, this ontology reuse should be approached in a
17 methodological way. The Ontological Resource Reuse
18 Process proposed by [32] describes the set of activities
19 to be performed for the reuse of existing ontological
20 resources.

- 21 1. Ontology Search. This activity consists in finding
22 appropriate ontological resources that meet the
23 requirements aimed to be satisfied. According
24 to [33], the existing ontology catalogues such as
25 LOV³ presented in [34] or LOV4IoT presented
26 in [35] (specialised in ontologies related to IoT)
27 can ease this task.
- 28 2. Ontology Assessment. This activity deals with
29 assessing the usability of an ontology with re-
30 spect to the requirements previously defined.
31 However, this may end up being a laborious
32 task due to the different criteria that may make
33 ontologies suitable for a certain use case. Fur-
34 thermore, the frequent scarce documentation of
35 ontologies may hinder this activity as potential
36 reusers may not understand the analysed ontol-
37 ogy.
- 38 3. Ontology Comparison. In this activity, assessed
39 ontologies should be compared according to cri-
40 teria that encompass the content of the ontol-
41 ogy, the organisation of these contents, the lan-
42 guage in which it is implemented, the method-
43 ology that has been followed to develop it, the
44 software tools used to build and edit the ontol-
45 ogy, and the costs of the ontology as suggested
46 by [36].

- 1 4. Ontology Selection. After assessing and compar-
2 ing ontologies, the most appropriate one or ones
3 have to be selected and reused by integrating
4 them in the new ontology being developed.

5 3.3. Benefits of ontology-based approaches

6 Ontology-based approaches are proved to bring
7 many advantages. Annotating raw data with terms
8 coming from ontologies not only allows a better rep-
9 resentation of the data itself, structuring it and set-
10 ting formal types, relations, properties and constraints
11 that hold among them, but also enables representing
12 data coming from multiple sources in a uniform way,
13 thereby supporting data integration [37] and data inter-
14 operability at a semantic level [38–41]. Furthermore,
15 additional background knowledge about a domain can
16 be added to the set of available data with ontologies.
17 This leads to the enrichment of the data set at hand,
18 as well as enabling the application of indexing tech-
19 niques to ensure the retrieval and navigation through
20 related resources [42]. Last but not least, after a seman-
21 tic annotation process, data is more domain-oriented
22 than the original source and allows more application-
23 independent solutions. Consequently, there is no need
24 for the user to be aware of the underlying structure of
25 the raw data.

30 4. Making Machine Learning Systems 31 Accountable

32 Towards the achievement of trustworthy AI sys-
33 tems, this article proposes an ontology-based ap-
34 proach aimed at providing Machine Learning systems
35 with accountability. The proposed approach is model-
36 agnostic, that is, it holds for any Machine Learning
37 method, so developers are free to use their preferred
38 Machine Learning model. This approach consists of
39 three phases as shown in Figure 1.

40 The first phase is related to the development of
41 the predictive model and its deployment in production
42 where it will be executed. In the second phase both
43 the procedure followed to develop the deployed predic-
44 tive model and the results produced by the predictive
45 model are annotated with the adequate ontology terms.
46 As for the third phase, it is responsible for managing
47 the annotations of the previous phase and facilitating
48 their exploitation by users. Next, each of these phases
49 is outlined.

50 ³<http://lov.linkeddata.es>

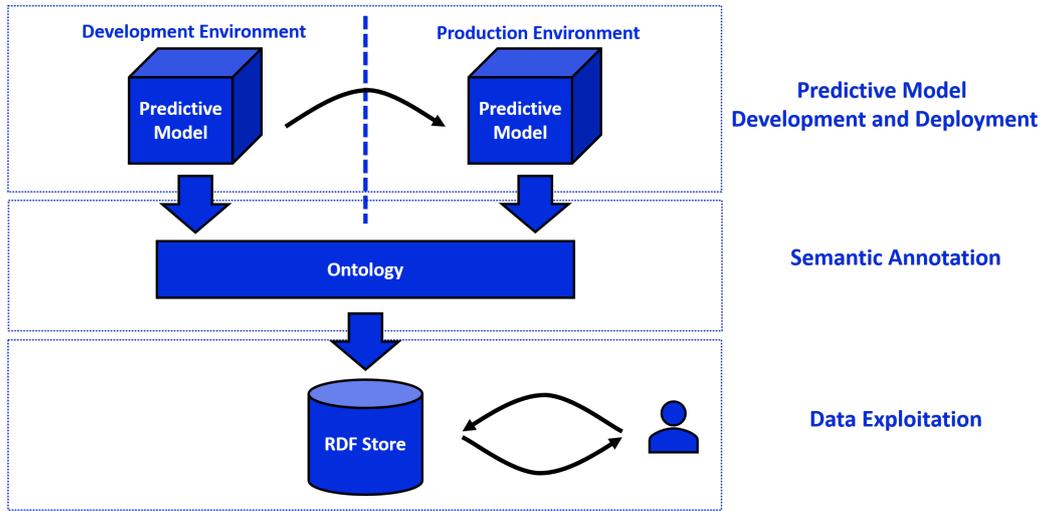


Fig. 1. Outline of the proposed ontology-based approach.

4.1. Predictive Model Development and Deployment phase

The initial phase consists in developing the predictive model that will identify the likelihood of future or unknown outcomes. To do so, the application of statistical algorithms and Machine Learning techniques to exploit historical data will be necessary.

Depending on the type of the problem addressed, the quality and the amount of data available, different algorithms may provide better results than others. Furthermore, the adequate fine-tuning of the hyperparameters of the algorithms may have a direct effect on the performance of the final model, and additionally, the performance of the model can be evaluated according to different metrics.

Once the models are developed, they need to be deployed, or in other words, a process for making them available in production environments has to be followed. It is only once the models are deployed to production that they start adding value, thus making the deployment a crucial step. Once they are deployed, they can be exploited to make the forecasts that are aimed by other software systems or end users.

4.2. Semantic Annotation phase

In the second phase of the proposed approach, the relevant information that enables making a Machine Learning system accountable has to be annotated with adequate ontology terms. To do so, the information to be represented needs to be identified first.

In this regard, two areas of knowledge are distinguished: the forecast made by the predictive model, and the procedure followed for making such a forecast. Likewise, the latter procedure-related information can be divided in the information that addresses the training data and the information concerning the predictive model itself.

The training data can be characterised both by its features and its quality. On the one hand, these features may include the amount of data used, the dependent and independent variables considered, and the statistical characteristics such as the variance, mean or median of the data. On the other, the quality of the training data determines whether it meets a standard set by the data scientist or not, and it can be measured in terms of its completeness, accuracy or conformity among others. As for the predictive model, information related to the algorithm used and its hyperparameters may be of interest, as well as the performance assessed in development time.

Once the information to be represented is identified, the adequate ontology or set of ontologies to annotate that information has to be selected. Following the ontology reuse best practice, this approach is based on the Ontological Resource Reuse Process [32]. For the search of relevant ontologies, four reputed sources have been consulted: LOV and LOV4IoT ontology catalogues, and Google Scholar⁴ and ScienceDirect⁵ research databases. Additionally, for the assessment,

⁴<https://scholar.google.com>

⁵<https://www.sciencedirect.com>

comparison and the final selection of the ontology to be reused, the following set of basic ontology quality criteria defined by [21] has been followed:

- Having an explicit license that specifies that they can be used and under which conditions.
- Having enough documentation to understand the ontology purpose, domain and fundamentals, and determine whether it describes this domain appropriately or not.
- Having a minimum metadata to help human users and computer applications understand the data as well as other important aspects that describes a data set.

This ontology reuse process has been followed for the two areas of knowledge to be represented: the forecasts themselves, and the procedures to make forecasts.

4.2.1. Representing forecasts

Currently, there are many ontologies which could be used for representing events or activities, the result of which is an estimate of the value of a quality of the feature of interest, obtained using a specific procedure. A thorough analysis of ontologies covering such a domain can be found in [21], and it can be concluded that the SOSA/SSN ontology proposed by [43, 44] may be one of the most appropriate ontologies for representing forecasts. However, SOSA/SSN ontology's admission of different models to represent the same state of affairs may derive in interoperability problems, which is why it was discarded for the approach proposed in this article. Instead, the EEP SA ontology⁶ proposed by [45] was selected to be reused, as it was developed on the basis that a proper axiomatisation shapes the set of admitted models better, and therefore, establishes the ground for a better interoperability.

Although being developed for supporting a data analyst assistant in energy efficiency and thermal comfort problems in buildings [46], the backbone of the EEP SA ontology is defined as a combination of three ODPs that can be used as basic building blocks to address similar problems in different domains. These ODPs try to be minimal in the number of classes and properties offered, but include appropriate ontology axioms that allow proper inferences. Namely, the three ODPs are the AffectedBy ODP⁷, the

Execution-Executor-Procedure (EEP) ODP⁸ and the Result-Context (RC) ODP⁹. Thanks to the great flexibility provided by this ODP-based ontology engineering modelling solution, the combination of these three ODPs have led to the development of ontologies covering other domains such as the agrifood as presented in [47].

The AffectedBy ODP defines two classes representing features of interest (*aff:FeatureOfInterest*) and their qualities (*aff:Quality*) and three object properties: *aff:belongsTo*, *aff:affectedBy* and *aff:influencedBy*. The *aff:belongsTo* object property supports the notion that every quality belongs to the feature of interest it is intrinsic to (i.e. a quality cannot belong to different features of interest), thus following the conceptualisation defined in the DOLCE upper level ontology proposed by [48]. The *aff:affectedBy* object property relates a quality with another quality that it affects, and the *aff:influencedBy* object property relates a quality with the feature of interest that it influences.

The EEP ODP imports the AffectedBy ODP and its two classes, and additionally, it defines three more classes: *eep:Execution*, *eep:Executor*, and *eep:Procedure*. An individual of *eep:Execution* is an event (e.g. a forecast) upon a quality of a feature of interest, produced by an agent by performing a procedure. As for an individual of *eep:Executor*, it is an agent capable of performing tasks by following procedures. Lastly, an individual of *eep:Procedure* describes the workflow, protocol, plan, algorithm, or computational method to be executed by agents to produce an event. Furthermore, the *eep:madeBy*, *eep:usedProcedure*, and *eep:onQuality* object properties are introduced in the EEP ODP. The *eep:madeBy* object property links an execution to the agent that performs the action, the *eep:usedProcedure* object property links an execution to the procedure that describes the task to be performed; and the *eep:onQuality* object property links an execution to the quality concerned by the execution. These three functional object properties, combined with a set of property chain axioms defined in the EEP ODP, allow the inference of the remaining object properties *eep:implements* linking executors to procedures, *eep:hasFeatureOfInterest* linking executions to features of interest, *eep:forQuality* linking executors to qualities, and *eep:forFeatureOfInterest* linking executors to features of interest.

⁶<http://w3id.org/eep sa>

⁷<https://w3id.org/affectedBy>

⁸<https://w3id.org/eep>

⁹<https://w3id.org/rc>

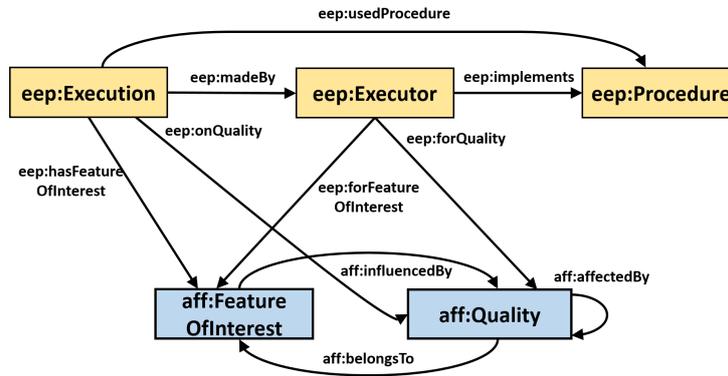


Fig. 2. The main classes and properties of the AffectedBy, EEP and RC ODPs.

The RC ODP aims at representing the results of the executions defined in the EEP ODP as well as their contexts. These results can be complex objects that usually include units of measurement, the measurement value, and some other optional parameters, but sometimes, a simple representation with a literal type value may suffice. Both complex and simple results can be modelled with the *rc:hasResult* object property and the *rc:hasSimpleResult* datatype property respectively. Furthermore, temporal and spatial aspects of a result are represented in the RC ODP with the *rc:hasGenerationTime* data property and the *rc:hasTemporalContext* object property for the former, and the *rc:hasSpatialContext* object property for the latter. Figure 2 shows the main classes and relationships defined in the three EEP core ODPs.

These three ODPs are published in the ODP repository [OntologyDesignPatterns.org](http://ontologydesignpatterns.org)¹⁰ and they are available online with a CC BY 4.0 license. They have a well-presented documentation, careful metadata with explanatory descriptions of the intended meanings of their terms, and alignments to other domain ontologies such as the SOSA/SSN ontology or W3C's PROV-O ontology¹¹ to ensure clarity in modelling and avoid errors that may have unintended reasoning implications [49].

4.2.2. Representing predictive procedures

The other area of knowledge that needs to be represented with adequate ontological terms is the one concerning the predictive procedures used for achieving forecasts. The existing ontologies in this domain are not as abundant as for the previous one, although there

are still ontologies covering Machine Learning experiments and different areas of the data mining, such as the OntoDM-core ontology described in [50] or the DMOP ontology presented in [51]. However, there is a gap between these ontologies, which definitely hampers an ideal interoperable scenario. Towards reducing such a gap and achieving a higher level of interoperability among those resources, a global standard schema is needed: the ML-Schema.

The ML-Schema¹² developed within the W3C Machine Learning Schema Community Group¹³ and described in [52] is an ontology that provides a set of classes, properties, and restrictions to represent different aspects of Machine Learning processes. On the one hand, the resources to describe the data used as input and their characteristics and quality are offered. On the other, resources for describing the implementations, algorithms used to develop models and their hyperparameters are defined. Finally, the developed models, their characteristics and the evaluation obtained in the training phase can also be represented with the ontology. Figure 3 shows the main classes and relationships defined in the ML-Schema.

The ML-Schema is published in the LOV catalogue and it is available online with a W3C Community Contributor License Agreement. Even though it has a complete documentation page, the metadata associated to the resources described in the ontology are incomplete. The guidelines proposed by [53], which are considered one of the most complete ontology metadata guidelines to date, are not met in many cases. For example, there is no human readable label for any of the defined terms, and additionally, there are

¹⁰<http://ontologydesignpatterns.org/>

¹¹<https://www.w3.org/TR/prov-o/>

¹²<http://www.w3.org/ns/mls>

¹³<https://www.w3.org/community/ml-schema/>

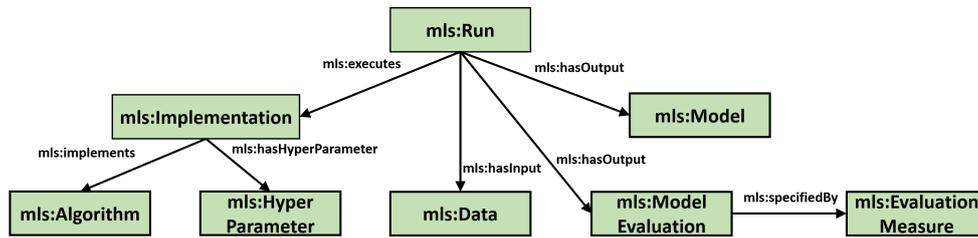


Fig. 3. The main classes and properties of the ML-Schema.

classes (e.g. *mls:DatasetCharacteristic*) and properties (e.g. *mls:hasValue*) which have no metadata associated, thus their intended meaning is not clear and their functionality is open to interpretation. Therefore, the W3C Machine Learning Schema Community Group should work on this issue to improve the vocabulary's reusability¹⁴. The ML-Schema is also mapped to more specific ontologies and vocabularies on Machine Learning such as the MEX vocabulary presented by [54].

Therefore, the approach presented in this article advocates the use of this set of ODPs and ontologies. However, since each scenario may have its own requirements, it may be possible that these ontological resources cannot satisfy them. In these cases, their extension is recommended, ideally, by reusing already existing ontological resources, thus following the best practices.

This semantic annotation phase could be performed with an ontology-driven editing framework to manually edit models and also semi-automatic tools to provide annotated data from data repositories could be used, such as platforms to map relational databases to RDF data, or data wrangling tools for more unstructured data. Regardless of the method used, the result will be a set of triples representing the target information.

4.3. Data Exploitation phase

Once the targeted data is semantically annotated with adequate ontological terms, the resulting triples must be stored in a repository where they remain accessible by the user responsible for checking the accountability of the corresponding Machine Learning system.

Triplestores or RDF stores are repositories that are optimised to handle this type of data, and they usually

provide SPARQL endpoints which are capable of receiving and processing information retrieval requests. On top of these SPARQL endpoints, additional services and interfaces can be developed in order to facilitate the interaction with the underlying SPARQL query language in which users might not be experts.

Furthermore, having the information structured in triples represented with proper ontology terms, contributes to the Linked Data concept. Therefore, this information could be enriched by interlinking it with other data sources and additional knowledge could be discovered [55–57].

5. Demonstrating the Ontology-Based Approach for Accountability

In order to illustrate the feasibility of the proposed approach, it has been implemented in a real-world scenario. In this scenario, a set of predictive models which forecast the electric demand of the next 24 hours for residential and small commercial buildings are developed and deployed into production. Namely, a total of 122 data-driven predictive models, one for each of the building units located in the island of Lanzarote, Spain. These models have to be executed on a daily basis as they are part of an energy-efficiency solution.

It is expected that these predictive models will be executed during a period of at least 1 year, therefore, a total of 44,530 forecasts (1 forecast per day per building unit x 122 building units x 365 days) are expected. Furthermore, under normal conditions, predictive models' performance degrade over time due to a change in the environment that violates the models assumptions [58], and models need to be retrained, which consists in re-running the process that generated the previous model on the new set of data available. However, since according to [59] the electric consumption patterns are strongly influenced by occupants' behaviour and lifestyle, a change in their habits (e.g. a family member leaving the household or a commercial build-

¹⁴An issue related to this matter is opened at the moment of writing this article in <https://github.com/ML-Schema/core/issues/25>.

ing extending their working hours) may make the typical retraining process insufficient and new predictive models may be necessary. This means that, each building unit, may need more than a single predictive model throughout this 1 year period considered.

These forecasts need to be accountable as they are part of an AI system for energy efficiency, and the managers of this system have expressed that the traditional approach of storing the necessary information in spreadsheets may not be the most good enough one, as the maintenance and handling of these files may be too complicated for the complex scenario described above. Additionally, the process of looking for some specific piece of information within these files is not straightforward and may end up being a tedious task.

Therefore, it is evident that managing the accountability in the scenario described is not a trivial task and that an approach supported by technologies that enable the management of the semantics and interrelationships of data, as well as the knowledge representation is essential. In this case, the ontology-based approach proposed in this article is followed.

5.1. Predictive Model Development and Deployment phase

All the predictive models needed are developed with the R programming language. They are all trained implementing the KNN algorithm of the caret¹⁵ package, and the value of the k hyperparameter is set by using the forward-chaining time series cross validation method. Furthermore, the forecasting effectiveness of the developed models is evaluated with the RMSE (Root Mean Squared Error).

The developed predictive models have been exported in the form of *.rds* files and put into production in an R Serve¹⁶ version 3.2.5 deployed in a Docker¹⁷ container. These predictive models are currently automatically executed once a day, using periodical tasks executed by a *cron daemon* process.

5.2. Semantic Annotation phase

The information to be annotated is retrieved from the R environment where the predictive models are generated and the execution of forecasts is performed. Furthermore, the semantic annotation of relevant data

with adequate ontological terms is automated to minimise potential performance issues and errors derived from manual practices. To do so, a service based on Apache Jena¹⁸, a Java framework for building Semantic Web and Linked Data applications, is developed.

For the sake of demonstrating the proposed ontology-based approach for accountability, let us consider the following simplified use case. A given predictive model was executed on 2020/11/25 at 07:00 and forecast that the building unit 02SX would have an electric consumption of 892 Wh on 2020/11/25 at 11:00. This predictive model was trained with data collected from the 02SX building unit between 2019/10/27 and 2020/03/15. This data set had an hourly frequency and a total of 29 data points were missing in this period of time due to connectivity issues of the sensor measuring the electric consumption. The features of the training set included, apart from the electric consumption, the hour when the measurement was made, the weekday and whether it was a working day or not. This predictive model was based on the R language's caret package's KNN algorithm implementation with the hyperparameter k set to 10, and obtained an RMSE of 242.03 Wh.

For making a semantic annotation of this use case, the set of ODPs and ontologies mentioned in the previous section are leveraged: the EEP SA ontology's three ODPs for representing the forecast, and the ML-Schema for representing the predictive procedure. In order to ease the understanding and readability of the proposed approach, the detailed explanation of the semantic annotation will be split into two: the forecast, and the predictive procedure. The triples representing this information can be found in Appendix A.

The forecast is defined as an instance of the *eep:Execution* class. It is made by (*eep:madeBy*) a given predictive model (*eep:Executor*) and produced by (*eep:usedProcedure*) following a given procedure represented as individual of the *eep:Procedure* class. The properties defined in the RC ODP are used for representing the actual value of the forecast (*rc:hasSimpleResult*), the instant when the forecast was generated (*rc:hasGenerationTime*) and the time for when the forecast is valid (*rc:hasTemporalContext*). Additionally, the forecast is related with the electric consumption of the building unit 02SX that predicts (represented as individual of the *aff:Quality* class) via the *eep:onQuality* object property, and this quality is

¹⁵<http://caret.r-forge.r-project.org/>

¹⁶<https://www.rforge.net/Rserve/>

¹⁷<https://www.docker.com/>

¹⁸<http://jena.apache.org/>

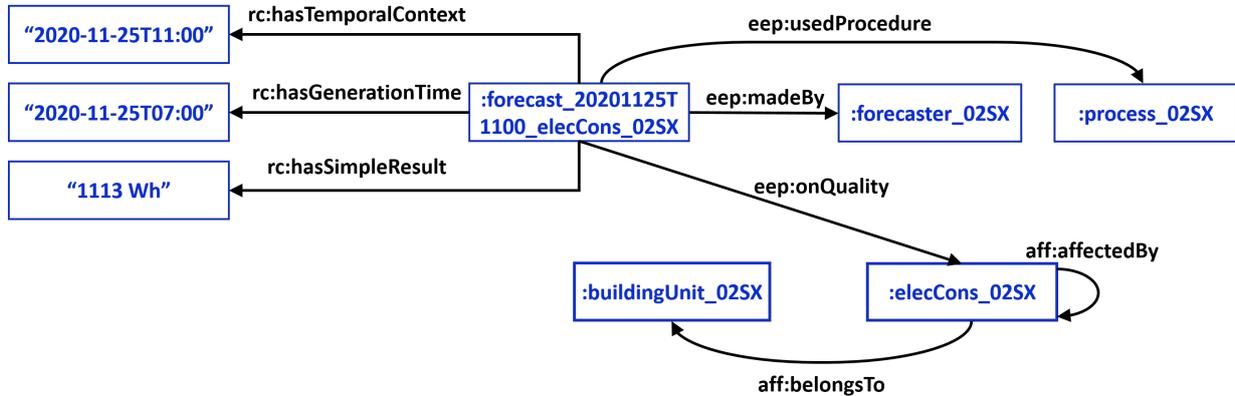


Fig. 4. Simplified graphic representation of the triples representing the 02SX building unit's electric consumption forecast.

linked with the the building unit 02SX (represented as an individual of the *aff:FeatureOfInterest* class) it belongs to. The triples describing the electric consumption forecast made for the 02SX are represented by Figure 4.

Regarding the procedure used for obtaining such a forecast, it is represented as an individual of the *mls:Process* class, which can be understood as a subclass of the broader *eep:Procedure* class. This procedure executes an R environment implementation (*mls:Implementation*) of the KNN algorithm (*mls:Algorithm*) with the *k* hyperparameter (*mls:Hyperparameter*) value set to 10. Additionally, the procedure is related to the data set used for the training process (*mls:Dataset*) via the *mls:hasInput* object property. This data set's features are represented with individuals of the *mls:DatasetCharacteristic* class, just like the data set's quality metrics via the *mls:hasQuality* object property. Finally, the resulting predictive model is represented as an individual of the *mls:Model* class and it is related with the process that generated it via the *mls:hasOutput* object property. Likewise, the model's evaluation (*mls:ModelEvaluation*) is specified by an individual of the *mls:EvaluationMeasure* class (in this case representing the RMSE) via the *mls:specifiedBy* object property and with a value of 242.03 Wh. The predictive model and its features are characterised by the triples represented in Figure 5.

5.3. Data Exploitation phase

The storage of the semantically annotated data is automated with a service based on Apache Jena. And

in this use case, an Openlink Virtuoso¹⁹ Universal Server version 07.20.3217 is used as repository, which is deployed in a Docker container. Virtuoso offers a SPARQL endpoint which enables retrieving the desired information via SPARQL queries. In order to facilitate the interaction with, set of predefined parameterisable SPARQL queries have been designed.

Let us consider that, for the demonstration scenario that we have proposed above, the forecast of electric consumption for a given house is largely inaccurate for some time. Therefore, it is reasonable to think that the person in charge of the management of the predictive models deployed may ask:

- Which is the performance obtained in the training of the model that forecasts the electric consumption of the building unit 02SX on 25th November 2020 at 11:00?

This information can be discovered by instantiating and running the parameterisable SPARQL query shown in Listing 1. Wild card *\$FORECAST_QUALITY* needs to be replaced with the forecast quality's URI (i.e. *:elecCons_02SX*), and wild card *\$FORECAST_TIME* with the date when the forecast is valid (i.e. *2020-11-25T11:00*). The results obtained are shown in Table 1.

```

PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX eep: <https://w3id.org/eep#>
PREFIX rc: <https://w3id.org/rc#>
PREFIX mls: <http://www.w3.org/ns/mls#>

```

¹⁹<https://virtuoso.openlinksw.com/>

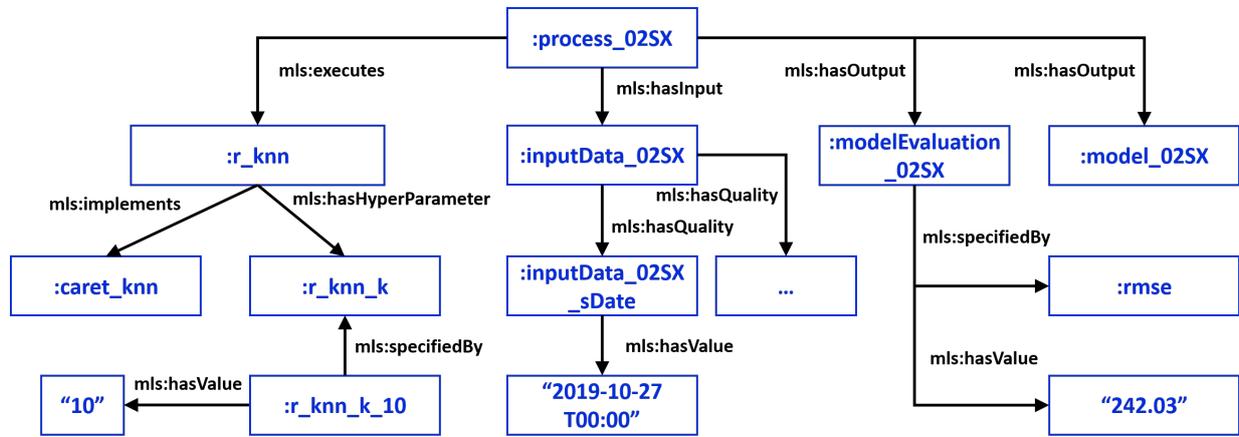


Fig. 5. Simplified graphic representation of the triples representing the example scenario's predictive model.

```

SELECT ?performanceMetric
      ?performanceValue
WHERE {
  ?forecast eep:onQuality
            $FORECAST_QUALITY;
  rc:hasTemporalContext $FORECAST_TIME;
  eep:usedProcedure ?procedure .

  ?procedure mls:hasOutput ?modelEval .

  ?modelEval rdf:type mls:ModelEvaluation;
  mls:specifiedBy ?performanceMetricURI;
  mls:hasValue ?performanceValue .

  ?performanceMetricURI rdfs:label
  ?performanceMetric .
}

```

Listing 1: SPARQL query for retrieving the performance obtained in the training of the model that forecast a certain quality for a certain instant of time.

Let us consider now that the person in charge of the management of the predictive models may wonder:

Table 1

Results obtained after running the SPARQL query shown in Listing 1, parameterised with the desired values.

?performanceMetric	?performanceValue
RMSE	243.03

- Which is the algorithm and hyperparameters used by the model that forecast the electric consumption of the building unit 02SX on 25th November 2020 at 11:00?

This information can be discovered by instantiating and running the parameterisable SPARQL query shown in Listing 2. Just like in Listing 1, wild cards *\$FORECAST_QUALITY* and *\$FORECAST_TIME* need to be replaced with the corresponding values. The results obtained are shown in Table 2.

```

PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX eep: <https://w3id.org/eep#>
PREFIX rc: <https://w3id.org/rc#>
PREFIX mls: <http://www.w3.org/ns/mls#>

SELECT ?algorithm ?hyperparameter
      ?hyperparamValue
WHERE {
  ?forecast eep:onQuality
            $FORECAST_QUALITY;
  rc:hasTemporalContext $FORECAST_TIME;
  eep:usedProcedure ?procedure .

  ?procedure mls:executes ?implementation .

  ?implementation mls:implements
  ?algorithmURI;
  mls:hasHyperParameter
  ?hyperparameterURI .

  ?hyperparameterSetting mls:specifiedBy
  ?hyperparameterURI;
  mls:hasValue ?hyperparamValue .

  ?algorithmURI rdfs:comment ?algorithm .

  ?hyperparameterURI rdfs:label
  ?hyperparameter .
}

```

Listing 2: SPARQL query for retrieving the algorithm and hyperparameters of the model that forecast a certain quality for a certain instant of time.

Table 2

Results obtained after running the SPARQL query shown in Listing 2, parameterised with the desired values.

?algorithm	?hyperparameter	?hyperparamValue
KNN implementation of caret R package	k	10

These two predefined parameterisable SPARQL queries are just illustrative examples that demonstrate the validity of the proposed approach for ensuring the accountability of Machine Learning systems. More SPARQL queries can be defined to retrieve the information sought by the end users.

6. Conclusions

The current adoption, deployment and application of AI systems is not as wide as it could be expected, mainly due to a lack of trust of users. Although XAI has been intensively studied lately, there are other factors that are equally important for achieving trustworthy AI systems. This article proposes an approach to address the accountability of Machine Learning systems based on the adequate representation of data, processes and workflows via ontologies.

The presented approach consists of three phases. The first one deals with the development of the predictive models and their deployment into a production environment where the actual forecasts are performed. The second phase annotates the relevant information coming from predictive models and forecasts with the adequate ontology terms. As for the third phase, it is in charge of storing the annotated data and providing the means to exploit it for ensuring the accountability of the targeted Machine Learning system.

After demonstrating the validity of the proposed approach in a real-world scenario, it can be concluded that the Semantic Technologies, thanks to their capabilities of managing the semantics and interrelationships of data as well as representing knowledge, are suitable for achieving the accountability of Machine Learning systems.

However, the author of this article suggests that the usage of Semantic Technologies for trustworthy AI systems should receive more attention from the Semantic Web community. As a matter of fact, the full potential of Semantic Technologies to fill existing gaps and unsolved challenges towards trustworthy AI systems is yet to be unlocked. This article is aimed at paving the way for future research in this direction.

6.1. Future Work

Although this article makes some contributions in the usage of ontologies towards the achievement of accountability of Machine Learning systems, there are many aspects that should be further researched.

The proposed approach should be tested for AI systems solving complex multiobjective problems. It is possible that the interaction between different components of such AI systems may require from additional ontological resources that are not covered by the ontologies currently considered in the presented approach. Additionally, other ontologies should be considered such as the Data Quality Vocabulary²⁰ [60] for describing the quality of the predictive models' training data.

Interaction with the proposed approach could be further improved with the design and implementation of a set of GUIs (Graphical User Interfaces), similar to how it has been done for [61]. This would help abstracting the users from the underlying SPARQL language in which they may not be experts.

Last but not least, as it has been previously concluded, the Semantic Technologies offer some benefits that could definitely contribute to the achievement of Trustworthy AI systems. Therefore, the research of the Semantic Technologies as a whole for solving other related factors such as fairness, explainability or transparency is also of interest.

Acknowledgements

This work is partly supported by the project 3KIA (KK-2020/00049), funded by the SPRI-Basque Government through the ELKARTEK program. Furthermore, the Protégé resource has been used, which is supported by grant GM10331601 from the National Institute of General Medical Sciences of the United States National Institutes of Health.

²⁰<https://www.w3.org/TR/vocab-dqv>

Appendix A. RDF examples

This appendix shows the RDF representation of the examples used in the article. For the sake of understandability the Turtle serialisation format has been used.

```

@prefix : <http://example.com/> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix aff: <https://w3id.org/affectedBy#> .
@prefix eep: <https://w3id.org/eep#> .
@prefix rc: <https://w3id.org/rc#> .
@prefix cdt: <http://w3id.org/lindt/custom_datatypes#> .

:buildingUnit_02SX rdf:type aff:FeatureOfInterest .
:elecCons_02SX rdf:type aff:Quality .
:forecast_20201125T1100_elecCons_02SX rdf:type eep:Execution ;
    rdfs:comment "Forecast made for 02SX building unit's electric consumption
    on 2020/11/25 at 11:00"^^xsd:String .
:forecaster_02SX rdf:type eep:Executor .
:process_02SX rdf:type eep:Procedure .

:elecCons_02SX aff:belongsTo :buildingUnit_02SX .
:forecast_20201125T1100_elecCons_02SX eep:onQuality :elecCons_02SX ;
    eep:madeBy :forecaster_02SX ;
    eep:usedProcedure :process_02SX ;
    rc:hasGenerationTime "2020-11-25T07:00"^^xsd:dateTime ;
    rc:hasTemporalContext "2020-11-25T11:00"^^xsd:dateTime ;
    rc:hasSimpleResult "1113 W.h"^^cdt:energy .

```

Listing 3: RDF representation of the example scenario's forecast.

```

@prefix : <http://example.com/> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix mls: <http://www.w3.org/ns/mls#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

:process_02SX rdf:type mls:Process ;
    mls:executes :r_knn ;
    mls:hasInput :inputData_02SX ;
    mls:hasInput :feature00_02SX ;
    mls:hasInput :feature01_02SX ;
    mls:hasInput :feature02_02SX ;
    mls:hasInput :feature03_02SX ;
    mls:hasOutput :model_02SX ;
    mls:hasOutput :modelEvaluation_02SX .

:r_knn rdf:type mls:Implementation ;
    mls:implements :caret_KNN ;
    mls:hasHyperParameter :r_knn_k .

:caret_KNN rdf:type mls:Algorithm ;
    rdfs:label "KNN"^^xsd:String ;

```

```

1      rdfs:comment "KNN implementation of caret R package"^^xsd:String .
2
3      :r_knn_k rdf:type mls:HyperParameter;
4          rdfs:label "k"^^xsd:String .
5
6      :r_knn_k_10 mls:specifiedBy :r_knn_k .
7      :r_knn_k_10 mls:hasValue "10"^^xsd:integer .
8
9      :inputData_02SX rdf:type mls:Dataset;
10         mls:hasQuality :inputData_02SX_sDate;
11         mls:hasQuality :inputData_02SX_eDate;
12         mls:hasQuality :inputData_02SX_NAs;
13         mls:hasQuality :inputData_02SX_Frequency .
14
15      :inputData_02SX_sDate mls:hasValue "2019-10-27T00:00"^^xsd:dateTime .
16      :inputData_02SX_eDate mls:hasValue "2020-03-15T00:00"^^xsd:dateTime .
17      :inputData_02SX_NAs mls:hasValue "29"^^xsd:integer .
18      :inputData_02SX_Frequency mls:hasValue "Hourly"^^xsd:String .
19
20      :feature00_02SX rdf:type mls:Feature;
21         rdfs:label "ElecCons"^^xsd:String;
22         rdfs:comment "The electric consumption value (response variable)"^^xsd:String .
23
24      :model_02SX rdf:type mls:Model;
25         rdfs:label "Model 02SX"^^xsd:String;
26         rdfs:comment "Model for the 02SX building unit"^^xsd:String .
27
28      :modelEvaluation_02SX rdf:type mls:ModelEvaluation;
29         mls:specifiedBy :rmse;
30         mls:hasValue "242.03"^^xsd:integer .
31
32      :rmse rdfs:label "RMSE"^^xsd:String;
33         rdfs:comment "Root-Mean-Square Error"^^xsd:String .

```

Listing 4: RDF representation of the example scenario's Predictive Model process.

References

- [1] S.S. ÓhÉigeartaigh, J. Whittlestone, Y. Liu, Y. Zeng and Z. Liu, Overcoming Barriers to Cross-cultural Cooperation in AI Ethics and Governance, *Philosophy & Technology* **33**, 571–593. doi:10.1007/s13347-020-00402-x.
- [2] M. Madsen and S. Gregor, Measuring human-computer trust, in: *11th Australasian conference on information systems (ACIS)*, Vol. 53, 2000, pp. 6–8.
- [3] J.-Y. Jian, A.M. Bisantz and C.G. Drury, Foundations for an empirically determined scale of trust in automated systems, *International journal of cognitive ergonomics* **4**(1) (2000), 53–71. doi:10.1207/S15327566IJCE0401_04.
- [4] B. Cahour and J.-F. Forzy, Does projection into use improve trust and exploration? An example with a cruise control system, *Safety science* **47**(9) (2009), 1260–1270. doi:10.1016/j.ssci.2009.03.015.
- [5] D. Gunning, Explainable artificial intelligence (xai), *Defense Advanced Research Projects Agency (DARPA), nd Web* **2**(2) (2017).
- [6] S.T. Mueller, R.R. Hoffman, W. Clancey, A. Emrey and G. Klein, Explanation in human-AI systems: A literature meta-review, synopsis of key ideas and publications, and bibliography for explainable AI, *arXiv preprint arXiv:1902.01876* (2019).
- [7] J. van der Waa, E. Nieuwburg, A. Cremers and M. Neerinx, Evaluating XAI: A comparison of rule-based and example-based explanations, *Artificial Intelligence* (2020), 103404. doi:10.1016/j.artint.2020.103404.
- [8] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti and D. Pedreschi, A survey of methods for explaining black box models, *ACM computing surveys (CSUR)* **51**(5) (2018), 1–42. doi:10.1145/3236009.
- [9] N. El Bekri, J. Kling and M.F. Huber, A study on trust in black box models and post-hoc explanations, in: *International Workshop on Soft Computing Models in Industrial and Environmental Applications*, Springer, 2019, pp. 35–46. doi:10.1007/978-3-030-20055-8_4.
- [10] J. Fox, The uncertain relationship between transparency and accountability, *Development in practice* **17**(4–5) (2007), 663–671. doi:10.1080/09614520701469955.
- [11] J.A. Kroll, S. Barocas, E.W. Felten, J.R. Reidenberg, D.G. Robinson and H. Yu, Accountable algorithms, *U. Pa. L. Rev.* **165** (2016), 633–705.
- [12] V. Beaudouin, I. Bloch, D. Bounie, S. Cléménçon, F. d’Alché-Buc, J. Eagan, W. Maxwell, P. Mozharovskiy and J. Parekh, Flexible and context-specific AI explainability: a multidisciplinary approach, *Available at SSRN 3559477* (2020). doi:10.2139/ssrn.3559477.
- [13] D. Oberle, How ontologies benefit enterprise applications **5**(6) (2014), 473–491. doi:10.3233/SW-130114.
- [14] A. Seeliger, M. Pfaff and H. Krcmar, Semantic Web Technologies for Explainable Machine Learning Models: A Literature Review, in: *Joint Proceedings of PROFILES 2019 and SEMEX 2019, 1st Workshop on Semantic Explainability (SemEx 2019), co-located with the 18th International Semantic Web Conference (ISWC ’19)*, PROFILES-SEMEX 2019, Vol. 2465, CEUR-WS, 2019, pp. 30–45. ISSN 1613-0073. http://ceur-ws.org/Vol-2465/semex_paper1.pdf.
- [15] R. Confalonieri, T. Weyde, T.R. Besold and F.M. del Prado Martín, TREPAN Reloaded: A Knowledge-driven Approach to Explaining Black-box Models **325** (2020), 2457–2464. doi:10.3233/FAIA200378.
- [16] S. Chari, O. Seneviratne, D.M. Gruen, M.A. Foreman, A.K. Das and D.L. McGuinness, Explanation Ontology: A Model of Explanations for User-Centered AI, in: *Lecture Notes in Computer Science*, Vol. 12507, Springer, 2020, pp. 228–243. doi:10.1007/978-3-030-62466-8_15.
- [17] C. Panigutti, A. Perotti and D. Pedreschi, Doctor XAI: an ontology-based approach to black-box sequential data classification explanations, in: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 2020, pp. 629–639. doi:10.1145/3351095.3372855.
- [18] F. Lecue, On the Role of Knowledge Graphs in Explainable AI, *Semantic Web* **11**(1) (2020), 41–51. doi:10.3233/SW-190374.
- [19] H.L.E.G.o.A.I. HLEG-AI, Ethics Guidelines for Trustworthy AI, 2019, last visited on 2021-02-08. <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>.
- [20] A.B. Arrieta, N. Díaz-Rodríguez, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins et al., Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI, *Information Fusion* **58** (2020), 82–115. doi:10.1016/j.inffus.2019.12.012.
- [21] I. Esnaola-Gonzalez, Semantic Technologies Towards Accountable Artificial Intelligence: A Poultry Chain Management Use Case (2020), 215–226. ISBN 978-3-030-63799-6. doi:10.1007/978-3-030-63799-6_17.
- [22] K.K. Breitman, M.A. Casanova and W. Truskowski, *Ontology in Computer Science*, in: *Semantic Web: Concepts, Technologies and Applications*, Springer London, London, 2007, pp. 17–34. ISBN 978-1-84628-710-7. doi:10.1007/978-1-84628-710-7_2.
- [23] R. Studer, V.R. Benjamins and D. Fensel, Knowledge engineering: principles and methods, *Data and knowledge engineering* **25**(1) (1998), 161–198. doi:10.1016/S0169-023X(97)00056-6.
- [24] O. Lassila and D. McGuinness, The role of frame-based representation on the semantic web, *Linköping Electronic Articles in Computer and Information Science* **6**(5) (2001), 2001.
- [25] N. Guarino, Semantic matching: Formal ontological distinctions for information organization, extraction, and integration, in: *International Summer School on Information Extraction*, Springer, 1997, pp. 139–170. doi:10.1007/3-540-63438-X_8.
- [26] M.C. Suárez-Figueroa, A. Gómez-Pérez and M. Fernández-López, *The NeOn Methodology for Ontology Engineering*, in: *Ontology Engineering in a Networked World*, M.C. Suárez-Figueroa, A. Gómez-Pérez, E. Motta and A. Gangemi, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 9–34. ISBN 978-3-642-24794-1. doi:10.1007/978-3-642-24794-1_2.
- [27] M. Poveda-Villalón, A. Fernández-Izquierdo and R. García-Castro, Linked Open Terms (LOT) Methodology, Zenodo, 2019. doi:10.5281/zenodo.2539305.
- [28] I. Esnaola-Gonzalez, J. Bermúdez, I. Fernandez and A. Arnaiz, *Two ODPs towards the notion of Influenceable Features of Interest*, in: *Advances in Pattern-based Ontology Engineering*, IOS Press, To-be-published.

- [29] A. Gangemi and V. Presutti, *Ontology Design Patterns*, in: *Handbook on Ontologies*, S. Staab and R. Studer, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 221–243. ISBN 978-3-540-92673-3. doi:10.1007/978-3-540-92673-3_10.
- [30] P. Hitzler, A. Gangemi and K. Janowicz, *Ontology Engineering with Ontology Design Patterns: Foundations and Applications*, Vol. 25, IOS Press, 2016.
- [31] E. Simperl, Reusing ontologies on the Semantic Web: A feasibility study, *Data & Knowledge Engineering* **68**(10) (2009), 905–925. doi:10.1016/j.datak.2009.02.002.
- [32] M. Fernández-López, M.C. Suárez-Figueroa and A. Gómez-Pérez, *Ontology Development by Reuse*, in: *Ontology Engineering in a Networked World*, M.C. Suárez-Figueroa, A. Gómez-Pérez, E. Motta and A. Gangemi, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 147–170. ISBN 978-3-642-24794-1. doi:10.1007/978-3-642-24794-1_7.
- [33] A. Gyrard, A. Zimmermann and A. Sheth, Building IoT based applications for Smart Cities: How can ontology catalogs help?, *IEEE Internet of Things Journal* **5** (2018), 3978–3990. doi:10.1109/JIOT.2018.2854278.
- [34] P.-Y. Vandenbussche, G.A. Atemezging, M. Poveda-Villalón and B. Vatant, Linked Open Vocabularies (LOV): a gateway to reusable semantic vocabularies on the Web, *Semantic Web* **8**(3) (2017), 437–452. doi:10.3233/SW-160213.
- [35] A. Gyrard, C. Bonnet, K. Boudaoud and M. Serrano, LOV4IoT: A second life for ontology-based domain knowledge to build Semantic Web of Things applications, in: *2016 IEEE 4th International Conference on Future Internet of Things and Cloud (FiCloud)*, IEEE, 2016, pp. 254–261. doi:10.1109/FiCloud.2016.44.
- [36] A. Lozano-Tello and A. Gómez-Pérez, Ontometric: A method to choose the appropriate ontology, *Journal of Database Management (JDM)* **15**(2) (2004), 1–18. doi:10.4018/jdm.2004040101.
- [37] N.F. Noy, Semantic integration: a survey of ontology-based approaches, *ACM Sigmod Record* **33**(4) (2004), 65–70. doi:10.1145/1041410.1041421.
- [38] S.N.A.U. Nambi, C. Sarkar, R.V. Prasad and A. Rahim, A unified semantic knowledge base for IoT, in: *2014 IEEE World Forum on Internet of Things (WF-IoT)*, 2014, pp. 575–580. doi:10.1109/WF-IoT.2014.6803232.
- [39] Y. Liao, M. Lezoche, H. Panetto, N. Boudjlida and E.R. Loures, Formal Semantic Annotations for Models Interoperability in a PLM environment, *IFAC Proceedings Volumes* **47**(3) (2014), 2382–2393, 19th IFAC World Congress. doi:10.3182/20140824-6-ZA-1003.02551.
- [40] Y. Liao, M. Lezoche, H. Panetto and N. Boudjlida, Semantic annotations for semantic interoperability in a product lifecycle management context, *International Journal of Production Research* **54**(18) (2016), 5534–5553. doi:10.1080/00207543.2016.1165875.
- [41] Y. Lin and H. Ding, Ontology-based Semantic Annotation for Semantic Interoperability of Process Models, in: *International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)*, Vol. 1, 2006, pp. 162–167. doi:10.1109/CIMCA.2005.1631259.
- [42] P. Andrews, I. Zaihrayeu and J. Pane, A classification of semantic annotation systems, *Semantic Web* **3**(3) (2012), 223–248. doi:10.3233/SW-2011-0056.
- [43] A. Haller, K. Janowicz, S. Cox, M. Lefrançois, K. Taylor, D.L. Phuoc, J. Lieberman, R. García-Castro, R. Atkinson and C. Stadler, The modular SSN ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation, *Semantic Web* **10**(1) (2019), 9–32. doi:10.3233/SW-180320.
- [44] K. Janowicz, A. Haller, S.J. Cox, D.L. Phuoc and M. Lefrançois, SOSA: A lightweight ontology for sensors, observations, samples, and actuators, *Journal of Web Semantics* **56** (2019), 1–10. doi:10.1016/j.websem.2018.06.003.
- [45] I. Esnaola-Gonzalez, J. Bermúdez, I. Fernandez and A. Arnaiz, EEPISA as a core ontology for energy efficiency and thermal comfort in buildings, *Applied Ontology* (To-be-published).
- [46] I. Esnaola-Gonzalez, J. Bermúdez, I. Fernandez and A. Arnaiz, Semantic prediction assistant approach applied to energy efficiency in Tertiary buildings, *Semantic Web* **9**(6) (2018), 735–762. doi:10.3233/SW-180296.
- [47] I. Esnaola-Gonzalez, I. Fernandez, E. García, S. Ferreiro, M. Gomez, I. Lázaro and A. García, Towards Animal Welfare in Poultry Farms through Semantic Technologies, in: *IoT Connected World & Semantic Interoperability Workshop (IoT-CWSI) 2019*, 2019. https://www.researchgate.net/publication/336848367_Towards_Animal_Welfare_in_Poultry_Farms_through_Semantic_Technologies.
- [48] S. Borgo and C. Masolo, Foundational choices in DOLCE, in: *Handbook on ontologies*, Springer, 2009, pp. 361–381. doi:10.1007/978-3-540-92673-3_16.
- [49] S. Cox, Ontology for observations and sampling features, with alignments to existing models, *Semantic Web* **8**(3) (2016), 453–470. doi:10.3233/SW-160214.
- [50] P. Panov, L. Soldatova and S. Džeroski, Ontology of core data mining entities, *Data Mining and Knowledge Discovery* **28**(5–6) (2014), 1222–1265. doi:10.1007/s10618-014-0363-0.
- [51] C.M. Keet, A. Ławrynowicz, C. d’Amato, A. Kalousis, P. Nguyen, R. Palma, R. Stevens and M. Hilario, The data mining optimization ontology, *Journal of web semantics* **32** (2015), 43–53. doi:10.1016/j.websem.2015.01.001.
- [52] G.C. Publio, D. Esteves, A. Ławrynowicz, P. Panov, L. Soldatova, T. Soru, J. Vanschoren and H. Zafar, ML-Schema: Exposing the Semantics of Machine Learning with Schemas and Ontologies, 2018.
- [53] D. Garijo and M. Poveda-Villalón, *Best Practices for Implementing FAIR Vocabularies and Ontologies on the Web*, in: *Applications and Practices in Ontology Design, Extraction, and Reasoning*, G. Cota, M. Daquino and G.L. Pozzato, eds, Studies on the Semantic Web, Vol. 49, IOS Press, 2020, pp. 39–54. ISBN 978-1-64368-142-9. doi:10.3233/SSW200034.
- [54] D. Esteves, D. Moussallem, C.B. Neto, T. Soru, R. Usbeck, M. Ackermann and J. Lehmann, MEX vocabulary: a lightweight interchange format for machine learning experiments, in: *SEMANTICS 15’ : Proceedings of the 11th International Conference on Semantic Systems*, 2015, pp. 169–176. doi:10.1145/2814864.2814883.
- [55] C. Gavankar, A. Kulkarni, Y.F. Li and G. Ramakrishnan, Enriching an academic knowledge base using Linked Open Data, in: *Proceedings of the Workshop on Speech and Language Processing Tools in Education*, 2012, pp. 51–60.

- 1 [56] L. Ruback, M.A. Casanova, A. Raffaetà, C. Renso and
2 V. Vidal, Enriching mobility data with Linked Open Data,
3 in: *Proceedings of the 20th International Database En-*
4 *gineering & Applications Symposium*, 2016, pp. 173–182.
5 doi:10.1145/2938503.2938550.
- 6 [57] K.F. Gracy, Enriching and enhancing moving images with
7 Linked Data, *Journal of Documentation* **74** (2018), 354–371.
8 doi:10.1108/JD-07-2017-0106.
- 9 [58] G. Widmer and M. Kubat, Learning in the presence of concept
10 drift and hidden contexts, *Machine learning* **23**(1) (1996), 69–
11 101. doi:10.1007/BF00116900.
- 12 [59] P. Lusi, K.R. Khalilpour, L. Andrew and A. Liebman, Short-
13 term residential load forecasting: Impact of calendar effects
14 and forecast granularity, *Applied Energy* **205** (2017), 654–669.
15 doi:10.1016/j.apenergy.2017.07.114.
- 16 [60] R. Albertoni and A. Issac, Introducing the Data Qual-
17 ity Vocabulary (DQV), *Semantic Web* **12**(1) (2021), 81–97.
18 doi:10.3233/SW-200382.
- 19 [61] I. Esnaola-Gonzalez, J. Bermúdez, I. Fernández and A. Arnaiz,
20 EROSO: Semantic Technologies Towards Thermal Comfort in
21 Workplaces, in: *European Knowledge Acquisition Workshop*,
22 Springer, 2018, pp. 519–533. doi:10.1007/978-3-030-03667-
23 6_33.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
511
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51