

Leveraging Ontologies to Document Bias in Data

Mayra Russo^{1,*}, Maria-Esther Vidal^{1,2}

¹L3S Research Center & Leibniz University of Hannover, Hannover, Germany

²TIB Leibniz Information Center for Science and Technology, Hannover, Germany

Abstract

Machine Learning (ML) systems are capable of reproducing and often amplifying undesired biases. This puts emphasis on the importance of operating under practices that enable the study and understanding of the intrinsic characteristics of ML pipelines, prompting the emergence of documentation frameworks with the idea that “any remedy for bias starts with awareness of its existence”. However, a resource that can formally describe these pipelines in terms of biases detected is still amiss. To fill this gap, we present the Doc-BiasO ontology, a resource that aims to create an integrated vocabulary of biases defined in the *fair-ML* literature and their measures, as well as to incorporate relevant terminology and the relationships between them. Overseeing ontology engineering best practices, we re-use existing vocabulary on machine learning and AI, to foster knowledge sharing and interoperability between the actors concerned with its research, development, regulation, among others. Overall, our main objective is to contribute towards clarifying existing terminology on bias research as it rapidly expands to all areas of AI and to improve the interpretation of bias in data and downstream impact.

Keywords

Bias, Ontology, Machine Learning, Trustworthy AI

1. Introduction

The breakthroughs and benefits attributed to big data and, consequently, to machine learning (ML) - or AI- systems [1, 2], have also resulted in making prevalent how these systems are capable of producing unexpected, biased, and in some cases, undesirable output [3, 4, 5]. Seminal work on bias (i.e., prejudice for, or against one person, or group, especially in a way considered to be unfair) in the context of ML systems demonstrates how facial recognition tools and popular search engines can exacerbate demographic disparities, worsening the marginalization of minorities at the individual and group level [6, 7]. Further, biases in news recommenders and social media feeds actively play a role in conditioning and manipulating people’s behavior and amplifying individual and public opinion polarization [8, 9]. In this context, the last few years have seen the consolidation of the *Trustworthy AI* framework, led in large part by regulatory bodies [10], with the objective of guiding commercial AI development to proactively account for ethical, legal, and technical dimensions [11]. Furthermore, this framework is also accompanied by the call to establish standards across the field in order to ensure AI systems are safe, secure and fair upon deployment [11]. In terms of AI bias, many efforts have been concentrated in devising methods that can improve its identification, understanding, measurement, and

AEQUITAS 2024: Workshop on Fairness and Bias in AI | co-located with ECAI 2024, Santiago de Compostela, Spain

*Corresponding author.

✉ mrusso@l3s.de (M. Russo); maria.vidal@tib.eu (M. Vidal)

🆔 0000-0001-7080-6331 (M. Russo); 0000-0003-1160-8727 (M. Vidal)

© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



mitigation [12]. For example, the special publication prepared by the National Institute of Standards and Technology (NIST) proposes a thorough, however not exhaustive, categorization of different types of bias in AI beyond common computational definitions (see Figure 1 for core hierarchy) [13]. In this same direction, some scholars advocate for practices that account for the characteristics of ML pipelines (i.e., datasets, ML algorithms, and user interaction loop) [14] to enable actors concerned with its research, development, regulation, and use, to inspect all the actions performed across the engineering process, with the objective to increase trust placed not only on the development processes, but on the systems themselves [15, 16, 17, 18].

In addition to human-readable (i.e., textual descriptions in a format that humans can read and understand) documentation frameworks for machine learning pipelines [19, 15, 20], semantic data models (e.g., ontologies, knowledge graphs) can also play a crucial role in enhancing the accuracy and interpretability of ML systems [21], as well as to perform “bias assessment, representation, and mitigation” tasks [22], in a way that is also machine-readable (i.e., makes available a fine-grained description of data in a format manageable by computers). This characteristic improves the findability, accessibility, interoperability and reusability (*FAIR*) of data-centric resources in the Web [23, 24]. Ontologies to model existing ML fairness metrics [25, 26], as well as the semantic specifications to catalog risks in terms of compliance and conformance of AI systems under the EU’s AI Act¹ [27, 28] have been proposed, however, a resource that can formally describe ML pipelines, and provides a vocabulary to characterize them in terms of measured biases is still amiss.

Proposed Solution We propose an ontology-driven approach to describe and document biases detected across machine learning pipelines. Here, we refer to documentation as the process of generating metadata represented in formats understandable by humans and also by machines [29], where formal data models like ontologies and controlled vocabularies provide standardized concepts for expressing this metadata. Our ontology, *Doc-BiasO*, is a resource developed with the objective to introduce an integrated vocabulary system of AI-related biases as defined in the literature and their measures; represent their relationships with other relevant terminology, i.e., datasets, ML systems, fairness, harm, risk; and semantically annotate ML pipelines based on bias measures values. The version presented here has 389 classes, 72 object properties, and 28 data properties.

Contributions: Concisely, our contributions are the following:

1. *Doc-BiasO*, an integrated vocabulary system of ML-related biases;
2. an ontology-based approach to document bias in ML pipelines;
3. a technical evaluation of *Doc-BiasO*.

The remainder of this paper is structured as follows: Section 2 introduces relevant Semantic Web concepts and presents a review of related literature. Section 3, details the design of *Doc-BiasO*. The results of the evaluation are reported in Section 4. Finally, Section 5 outlines our conclusions and future lines of work.

¹Annex III, European Council position

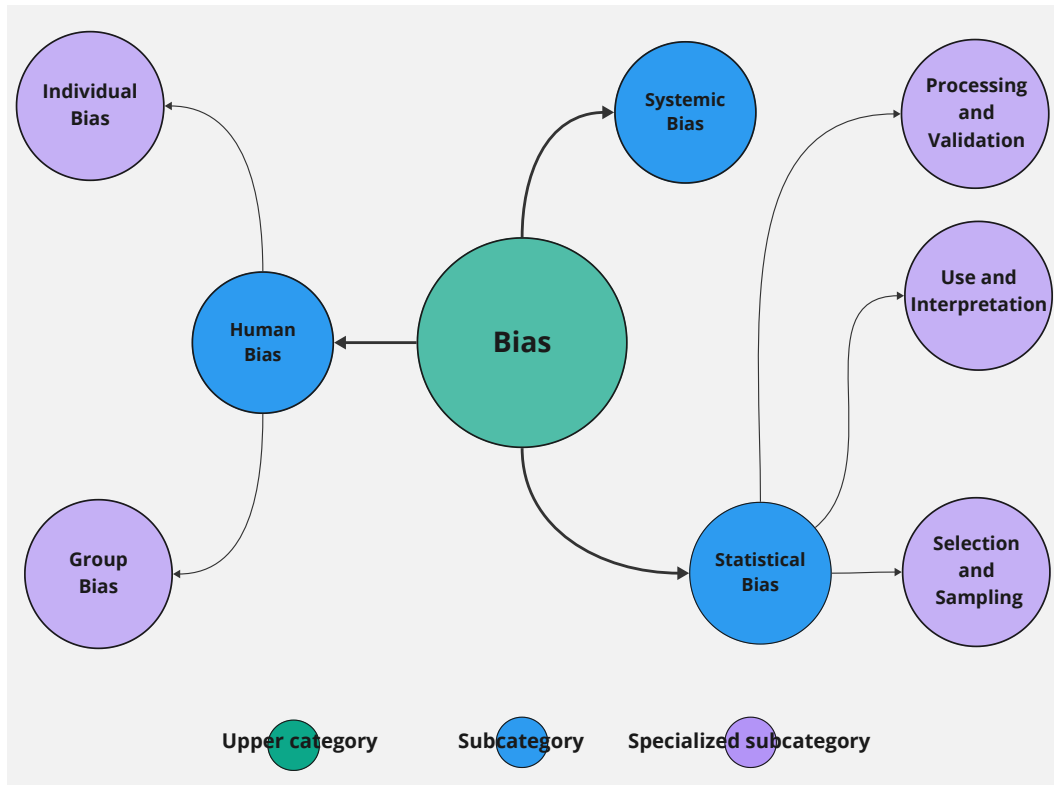


Figure 1: Types of Bias. Core categories of bias in relation to AI systems as per the NIST [13]. In the center and in turquoise, we depict the biggest circle for the Bias concept, as the most abstract and highest category. Via thicker arrows, we depict three smaller blue circles representing three main sub-categories of Bias. Using finer arrows that draw out from the three sub-categories, we depict smaller soft violet circles, to represent further sub-categories in the bias hierarchy.

2. Background and Related Work

Ontologies and Machine Learning Gruber [30] defines an ontology as a formal, explicit specification of a shared conceptualization that is characterized by high semantic expressiveness required for increased complexity. Ontologies include abstract concepts or classes, represented as nodes, and predicates representing the relations of these classes, edges in an ontology; the meaning of the predicates is represented using rules. Ontologies are specified using knowledge representation models, making the expressiveness of the ontology dependent on the expressive power of the representation model. The Resource Description Framework (RDF)² enables the description of entities in terms of classes and properties; while subsumption relations between classes and properties can be modelled with the RDF Schema (RDFS).³ More expressive formalisms like the Ontology Web Language (OWL)⁴ make available a larger number of operators

²<https://www.w3.org/RDF/>

³<https://www.w3.org/TR/rdf12-schema/>

⁴<https://www.w3.org/OWL/>

which enable the representation not only of classes, properties, and subsumption relations, but also class and property constraints, general equivalence relations, and restrictions of cardinality. Several examples of the usefulness of context-aware ontologies for bias awareness and mitigation in ML systems are explored in the work presented in [22].

In the context of bias modelling, the **Bias Ontology Design Pattern (BODP)** [31] is one of the first works to propose a formalization for the bias concept. Its objective is to capture a high-level representation of bias as an abstract term and not necessarily in the context of ML systems. We re-use part of BODP as a building block to repurpose it for scope and intended use of Doc-BiasO, which is to document bias in AI pipelines. Similar to our work, the **fairness metrics ontology (FMO)** [25, 26] models fairness metrics (`fmo:fairness_metric`) from the literature and relates them to their use-case. The conceptualization of bias and fairness in relation to ML systems are often intertwined; however, distinctions between both concepts need to be made explicit, as they are not always used in conjunction, nor to study the same phenomena [32, 33]. Fairness, in relation to ML (*fair-ML*), takes the form of algorithmic interventions that incorporate mathematical formalizations of moral or legal notions for the fair treatment of different populations into ML pipelines. These interventions aim to prompt ML models to satisfy statistical non-discrimination criterion for a given subpopulation [4]. In our case, we focus on modelling biases in data identified in the literature and the existing measures defined to detect them. Specifically, we propose a descriptive bias vocabulary that can be used and incorporated into varying frameworks as needed and that can be extended to further semi-automatize documentation tasks. Concepts and relations pertaining to bias are not made explicit in the current version of FMO, however, we consider both ontologies to be complementary, thus we re-use FMO to foster the development of a comprehensive vocabulary that provides coverage of terminology that pertains to the responsible development of ML systems. We follow a similar approach with the **AI Risk Ontology (AIRO)** [27], and by-effect, the **Vocabulary of AI Risks (VAIR)** [28]. In this case, risk in relation to ML systems, under the broader label of AI, is defined as systems that are likely to cause serious harms to health, safety, or fundamental rights of individuals as per European Union (EU) Law. These works are ontology-driven approaches to account for the compliance and conformance of AI systems under the EU’s AI Act’s specifications.⁵ Specifically, AIRO is a modular ontology created to identify whether an AI system is classified as high-risk, whilst VAIR provides semantic specifications for cataloging AI risks, re-using core concepts in AIRO (e.g., `airo#Risk`, `airo#Consequence`). Lastly, [34] proposes a descriptive framework (**ACROCPoLis**) to describe ML systems and their societal impact by making explicit the interrelations and diverging perspectives of relevant stakeholders (individuals, groups of people, institutions). While this is beyond the scope of our work, should the conceptual model be formalized and made publically available, a study for re-use and extension of Doc-BiasO ontology would be undertaken in a future iteration.

The Semantic Web community has also proposed other technical solutions to improve the interpretability and transparency of machine learning pipelines. The provenance ontology, PROV-O [35], enables the representation of provenance information generated by different entities, and can be easily applied to multiple contexts (i.e., training datasets). Standard schemas for data mining and machine learning algorithms, such as the Machine Learning Schema

⁵Annex III, European Council position

Table 1
Core Concepts in Doc-BiasO.

Concept	Definition
Bias	A concentration on, or interest in one particular area or subject. Whilst a more value-laden definition, conceptualizes bias as prejudice for, or against one person, or group, especially in a way considered to be unfair [45].
Application	The use, purpose or application of a machine learning system. Examples include, recommenders, speech recognition, etc.
ML Task	Task or ML Problem is the formal description of a process that needs to be completed (e.g., based on inputs and outputs) [36].
Dataset	A collection of data, published or curated by a single source, and available for access or download in one or more representations [36].
Harm	Adverse lived experiences resulting from an ML system’s deployment and operation in the world [46, 47].
Bias Measure	A quantitative metric or indicator that assesses the presence and extent of bias in a particular context, via predefined thresholds [48].

(MLS) ontology [36], and the Description of a Model (DOAM)⁶ ontology, provide fine-grained vocabularies to represent ML models characteristics. Moreover, the issue of reproducibility in ML has also been addressed [37]. Correspondingly, the Data Catalog Vocabulary (DCAT) [38] enables the fine-grained description of datasets and data services in a catalog using a controlled and rich vocabulary. Adhering to ontology engineering best practices [30], all these ontologies and vocabularies have been re-used in the composition of Doc-BiasO.

Documentation Frameworks and Machine Learning The opaqueness of the inner processes of ML systems can hinder the understanding of how they work. [19], Gebru et al. [15], [20], Hupont et al. [39], thus advocate for the production of value oriented, human-readable documentation for datasets (Data Statements for Natural Language Processing, Datasheets for Datasets), ML models (Model Cards for Model Reporting and Use Case Cards). Doc-BiasO aims to follow their stride by combining the different components of the AI pipeline (input, model, output data) to produce comprehensive descriptions in human- and machine-readable format of data-driven pipelines. Other documentation approaches, such as Sun et al. [40] introduce a tool to assess fitness for use of datasets. This automated data exploration tool delimits its focus to three dimensions: representativeness, bias, and correctness. In a similar line, [41] introduces a bias visualization tool for computer vision datasets. This exploration tool narrows down their assessment to three sets of metrics: object-based, gender-based and geography-based dimensions. Further, interactive tools– developed by industries– (e.g., [42, 43, 44]) enable dataset exploration, visualization, and comparison. The extensible and modular design of Doc-BiasO, allows users to describe and document their data-driven pipelines, and seamlessly incorporates additional descriptive dimensions and components as needed. Further, the underlying knowledge-driven framework prompts the integration and fine-grained description of multiple data sources, and leverages reasoning capabilities for enhanced data analytics.

⁶<https://www.openriskmanual.org/ns/doam/index-en.html>

3. Design and Implementation

In this section, we describe the design stages of Doc-BiasO. We also describe its implementation and include an example of an instance.

3.1. Scoping out the Coverage of Doc-BiasO

To determine the scope of our ontology, we perform a domain and content analysis following a hybrid strategy. On the one hand, through our own position within a research project on bias in relation to the development and regulation of AI⁷, we have held fruitful discussions with experts researching different dimensions of bias from a multidisciplinary and critical point of view, i.e., [49, 12]. Further, these discussions have helped identify what concepts make up our universe of discourse, for instance, bias, ML model, dataset, task, application, fairness, harms, risks; as well as how these concepts interact or relate to each other. In Table 1, we have summarized the core concepts defined in our ontology. Each of these concepts represents the top-most abstract concept in a hierarchy of terms, with less abstract or more concrete concepts being defined as the ontology grows to give a broader coverage. For example, `Bias` is the most abstract representation, while `Representation Bias` is a more concrete type of bias.

The exchanges with researchers have also helped deepen our understanding and characterization of bias in data from a critical stance (e.g., there is never just one bias, bias detection is contextual, bias detection can depend on data modality, biases cannot be eradicated [12]) and to identify challenges not only in modelling bias, but also in relation to the underlying documentation process, primarily on how it should not be fully automatized. In developing a tool like our ontology, it is important to aim for a careful balance between an effective, useful and comprehensive vocabulary that supports streamlining documentation tasks, while at the same time, avoid dissuading practitioners from critical thinking when engaging in both documentation and bias analysis. The aim of both of these practices is to mitigate negative consequences arising from the deployment of ML systems. However, it is always possible that unintentionally through enforcing standardization or automation on practitioners, new gateways are created that worsen the problem. Some influencing factors are the lack of experience, domain knowledge, or the right incentives [46, 50, 51]. Ultimately, this rapport informs our design choices across all iterations of ontology engineering, makes us aware of the limitations of our technical tool, and creates opportunities for refinement in later versions.

On the other hand, the scope of our ontology is also informed by the growing body of literature on our topic of interest. In this case, we particularly rely on official reports, as is the NIST Special Publication 1270 [13], and by periodically identifying relevant work in order to gather background information for a rich vocabulary of biases- (e.g., [52, 33, 53, 14, 54]), while also consider emerging work on this topic published at venues such as: ACM FAccT⁸, AAAI/ACM AIES⁹, ACM EAAMO.¹⁰ Concisely, we pay attention to discerning bias, and its detection measures from fairness notions and their measures, combining keywords such as

⁷<https://nobias-project.eu/>

⁸<https://facctconference.org/>

⁹<https://www.aies-conference.com/>

¹⁰<https://eaamo.org/>

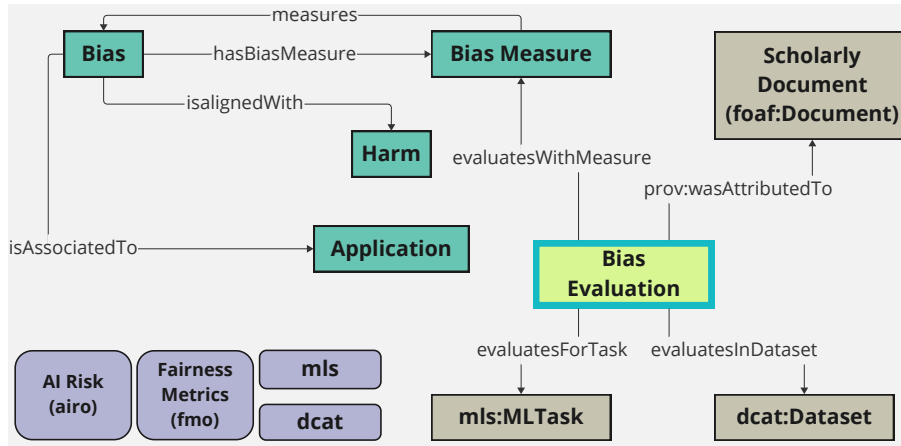


Figure 2: Conceptualization of the Doc-BiasO Ontology. Core concepts in the ontology are represented as classes, in color-coded boxes, to account for originating vocabularies. While object properties are drawn as directed arrows between classes. In purple colored boxes, relevant and prominently re-used vocabularies implemented in the representation of the universe of discourse.

“machine learning”, “artificial intelligence”, “datasets”, “bias”, “metrics”, and “bias mitigation”.

3.2. Doc-BiasO Design

To design and model our ontology, we adhere to ontology engineering best practices [30, 55]. As such, after the scope is determined and competency questions are defined, re-usable ontologies are identified following a layered approach (i.e., a foundational layer for general metadata and provenance, a domain-dependent layer to cover standards for the relevant area of use, a domain-dependent layer of ontologies specific to our problem of interest) [55].

We first specify the competency questions that emerged during the analysis phase and that represent the intended use of Doc-BiasO: a tool that can be integrated into AI documentation frameworks and that can offer the vocabulary required to characterize these pipelines; ideally, a resource that informs AI practitioners or researchers on the ways in which bias interacts with other components in the AI pipeline, and as a controlled repository as they study the development of a new measure and wish to survey those that already exist.

We then lay the foundation of our ontology by re-using ontologies such as: the SKOS data model [56], the PROV data model (PROV-O) [35], and the Friend of a Friend (FOAF) vocabulary [57]. The next layer incorporates standard schemas for data mining and machine learning algorithms, such as the Machine Learning Schema (MLS) ontology [36]. This schema provides fine-grained descriptions to represent the characteristics and intricacies of ML models. Similarly, the Data Catalog Vocabulary (DCAT) [38] enables the fine-grained description of datasets and data services in a catalog using a controlled and rich vocabulary. By extension, the Data Quality Vocabulary (DQV) [58] provides a framework and vocabulary to assess the quality of a dataset, offering an extensive catalog of quality metrics. For our third layer, we look at previous work on bias, specifically the BODP [31] and the Artificial Intelligence

Ontology (AIO).¹¹ The class AIO:Bias is our starting point, which we organize in hierarchies via `rdfs:SubClassOf`, as per the AIO modelling, and in order to represent different kinds of bias identified in the literature i.e., representation bias, popularity bias, demography bias. We build on the pattern and ontology, however, it does not suffice to our modelling needs. For this reason, all missing concepts are incorporated manually, as we set out to capture and explicitly document otherwise unstated assumptions about bias in relation to ML systems [59]. Critical data studies [47, 59] maintain that for bias detection tasks to be meaningful, practitioners must *reflect* on possible harms that can emerge upon the deployment of an ML system in dynamic societal and cultural contexts. Here, we emphasize thus on both, the importance of assisting practitioners via the development of tools that streamline tasks that may be perceived as a burden [50], while avoid dissuading them from reflecting about harms that could emerge from deploying these systems. For that reason, in our modelling we align scoped biases with harms, with the objective to make explicit the articulation of otherwise alleged, unstated negative consequences attributed to ML systems. However, our expectation and recommendation, is that users will enrich the proposed vocabulary with the results derived of their own explorations, in a similar line as with AI incident databases. Furthermore, bias is not singular, and highly context dependent, meaning that most biases are studied and defined in association to a particular ML application. To represent both of these concepts, we model `bias:Harm` and `bias:Application`. The central concept in our ontology is `bias:BiasMeasure`. This class represents a measure defined in some `foaf:Document`, evaluated in a `dcat:Dataset` (that has some characteristics), and for a particular `mls:MLTask`. `bias:BiasEvaluation` is the class that represents the n-ary relationship between entities schematized in the extended entity relationship model completed at the start of the design phase. Figure 2 illustrates a conceptual overview of the core classes and relationships of Doc-BiasO.

Towards a Comprehensive Vocabulary for Trustworthy AI The Trustworthy AI framework requires a comprehensive formal vocabulary that unifies approaches and contemplates terminology and concepts of ML pipelines, and in broader terms AI, holistically. This type of resource can contribute to the generation of metadata that primes reproducibility and traceability of research results [23, 24], a known issue in ML research and development [60, 37]. Moreover, it can help achieve a certain degree of standarization for the area. Motivated by this, we perform an analysis of the FMO [25] and VAIR [28] ontologies, as to determine their characteristics and how they fit into our model. We also do this with the aim to achieve a good balance between ontology re-use and down the line overhead derived from doing so [55]. Key takeaways are:

1. FMO complements Doc-BiasO by giving coverage to existing fairness metrics used to evaluate ML systems. Specifically, metrics pertaining to machine learning problems of classification and regression;
2. VAIR captures a wider scope of AI system deployment to instill accountability on an AI provider (i.e., a party that places the system on the market) and thus capture specifications of risky applications of AI from a regulatory point of view;
3. both ontologies represent bias, however, with differing modelling objectives. FMO organizes `fmo#Bias` in a hierarchy with seven subclasses, two of these are used in relation to

¹¹<https://bioportal.bioontology.org/ontologies/AIO?p=summary>

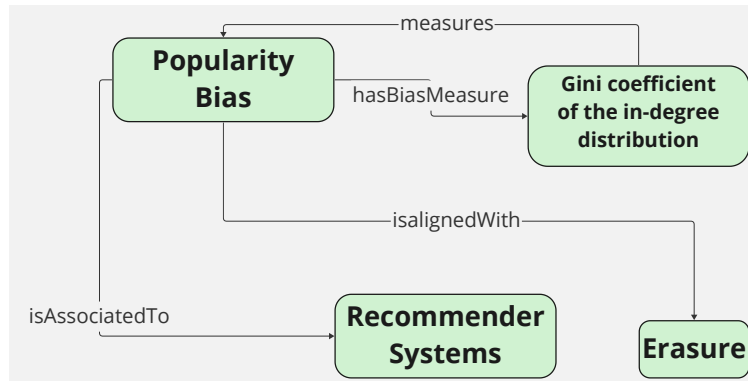


Figure 3: Conceptualization of an instance of Doc-BiasO. Instances of the Doc-BiasO ontology are represented with round-edge boxes and the color green. “Popularity Bias” is an instance of bias:Bias. Related classes are also exemplified.

some fairness metric. VAIR represents `vair#Bias` as a subclass of `airo#Consequence`.

To avoid constraining our modelling, we opt to not import either ontology in its entirety. When needed, we implement OWL axioms to assert class equivalence, i.e., `owl:equivalentClass`. Otherwise, we reference external concepts using annotation properties.

3.3. Doc-BiasO Specifications

Doc-BiasO Axiomatization The conceptualization of the Doc-BiasO ontology is specified using OWL logical axioms, given that OWL is formally defined in Description Logic. By using OWL to formalize our ontology, we enable consistency checks and logical inferences on a resulting RDF knowledge graph.

Further details can be found in the ontology documentation.¹²

Instantiating Doc-BiasO To showcase an instantiation of Doc-BiasO, we look at an example based on bias detection in relation to recommender systems, commonly implemented in online social networks.

The class Bias is instantiated as Popularity Bias. This bias is Associated With, an instance of the class Application, Recommender System and has a Bias Measure, "Gini coefficient of the in-degree distribution". In this example, Popularity Bias is Aligned With the instance of the class Harm, which is Erasure. We illustrate this in Figure 3.

4. Evaluation

4.1. Competency Questions

The domain analysis and scope definition of Doc-BiasO, as already described in Section 3.1, derived a set of competency questions that was also used to convey the requirements that

¹²<https://github.com/SDM-TIB/Doc-BIASO>

```

PREFIX skos: <http://www.w3.org/2004/02/skos/core#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX bias: <https://bias-project.x/bias/>

```

```

SELECT DISTINCT
  ?bias_1 (COUNT(DISTINCT ?biasMeasure_1) AS
  ?number_of_measures)
WHERE { ?bias_1 rdfs:subClassOf bias:Bias .
  ?biasMeasure_1 bias:measures ?bias_1}
GROUP BY ?bias_1

```

Listing 1: SPARQL Query for Competence Question Q4.1

would guide the engineering of our ontology. As part of the process, we tested and refined the Doc-BiasO ontology by implementing the formalization of the competency questions originally expressed in natural language as SPARQL queries.¹³ The queries were tested to make sure the results were the expected ones. The set of queries can be access through our GitHub.¹² To illustrate their adequacy, we continue with the example introduced earlier, and start by posing Q1 *”Given a particular bias, what is its definition?”*; our example uses Popularity Bias. Below the query result:

”When collaborative filtering recommenders emphasize popular items (those with more ratings) over other “long-tail”, less popular ones that may only be popular among small groups of users.”@en

This expected result is expressed as a `rdfs:Literal` in English. We follow this question by posing Q4.1 *”How many measures have been documented for it?”*. The results produced by executing the corresponding query, specified in Listing 1, are that for Popularity Bias, we have 3 measures. We choose the measure, Gini coefficient of the in-degree distribution, to learn more about it. We proceed to execute the query that corresponds to Q6. *what is its formalization?*. The corresponding SPARQL query is specified in Listing 2, with its execution projecting the definition for the chosen measure and the formalization for it in natural language.

As part of the evaluation process, we also report on the quality of Doc-BiasO. Table 2 summarizes the results obtained according to three indicators defined in [61].

4.2. Automatic Ontology Evaluation

This version of Doc-BiasO has also been validated with online tools to verify its consistency and syntactical validity, as well as to check for modelling anomalies or errors. First, we checked that our ontology is syntactically correct using the W3C RDF validation service.¹⁴ The results indicated a successful validation of our RDF document. Second, we checked for

¹³<https://www.w3.org/TR/sparql11-query/>

¹⁴<https://www.w3.org/RDF/Validator/>

```

PREFIX skos: <http://www.w3.org/2004/02/skos/core#>
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX bias: <https://bias-project.x/bias/>

```

```

SELECT DISTINCT
  ?biasMeasure_1 ?definition_1 ?formalization_1
WHERE {
  ?biasMeasure_1 rdfs:subClassOf bias:BiasMeasure ;
                skos:definition ?definition_1 ;
                bias:formalization ?formal_1
FILTER ( ( REGEX(str(?biasMeasure_1), "Gini", 'i')) ) }

```

Listing 2: SPARQL Query for Competence Question Q6

Table 2
Quality Indicators for Doc-BiasO.

Indicator	Results
Completeness	
Bias	All 51 subclasses have verifiable definitions based on the NIST report, $\frac{59}{51} = 115\%$.
Bias Measures	8 subclasses with verifiable definitions based on ongoing literature review, 24 instances based on 3 case studies.
Interoperability	
Using external vocabulary	$\frac{316}{389} = 81\%$
Used proprietary vocab	$\frac{73}{389} = 19\%$
Accessibility	http://ontology.tib.eu/DocBIASO/visualization

logical consistency by running the DL reasoning engine Pellet (v.2.2.0), as a plug-in for the Protégé open-source platform (v.5.6.1).¹⁵ We choose this engine as it is a complete reasoner. The results determined that Doc-BiasO is logically coherent and consistent. Finally, we scanned our ontology with the “OOPS! Ontology Pitfall Scanner” [62] to automatically dismiss the existence of modelling pitfalls; the evaluation results were also positive, as there were no bad practices detected by the tool.

5. Conclusions and Future Work

In this work, we presented Doc-BiasO, an ontology for bias measures found in the literature that can support the elaboration of documentation of bias in machine learning pipelines. Our objective is to contribute towards improving the interpretation of these pipelines in terms of biases captured, and the derived harms attributed to ML systems. Further, we make a call for a

¹⁵<https://protege.stanford.edu/software.php>

unified controlled vocabulary for the Trustworthy AI framework, and assess existing relevant work. We technically evaluated Doc-BiasO and showcase an example of its instantiation. Notwithstanding, our work is not without limitations. First, research on bias in ML, and by extension AI, is a fast-moving field, thus providing adequate and updated coverage with our tool is a challenge. Second, bias evaluation are highly complex and context dependent tasks. This means that our modelling cannot account for all potential existing biases, and that in general, bias analysis cannot be fully automated, requiring a human-in-the-loop. Third, our resources are yet to be evaluated by AI practitioners outside a research environment. Nevertheless, the addressed limitations are an opportunity for future work. In particular, we intend to add and expand on aspects left unmodeled in this version, and we will liaise with AI practitioners to evaluate the suitability of our tool in real world scenarios. We will also continue the development of a controlled vocabulary for Trustworthy AI, as this resource can foster effective communication between the different actors involved across the AI pipeline.

Acknowledgments

We thank Guillermo Climent-Gargallo, Sammy Sawischa and Yukti Sharma for their support during this research. Mayra Russo is supported by EU-Horizon 2020 research and innovation programme under the MCSA-grant agreement No. 860630, project: NoBIAS. Maria-Esther Vidal is partially supported by Leibniz Association, program "Leibniz Best Minds: Programme for Women Professors", project TrustKG-Transforming Data in Trustable Insights; Grant P99/2020.

References

- [1] H. Suresh, J. Guttag, A framework for understanding sources of harm throughout the machine learning life cycle, in: *Equity and Access in Algorithms, Mechanisms, and Optimization*, EAAMO '21, 2021. URL: <https://doi.org/10.1145/3465416.3483305>. doi:10.1145/3465416.3483305.
- [2] J. Riley, The elusive promise of ai: A second look, *Ubiquity 2021* (2021). URL: <https://doi.org/10.1145/3458742>. doi:10.1145/3458742.
- [3] R. A. Baeza-Yates, Big data or right data?, in: *Alberto Mendelzon Workshop on Foundations of Data Management*, 2013. URL: <https://api.semanticscholar.org/CorpusID:12577033>.
- [4] S. Barocas, M. Hardt, A. Narayanan, *Fairness and Machine Learning*, fairmlbook.org, 2019. <http://www.fairmlbook.org>.
- [5] S. Barocas, A. D. Selbst, Big data's disparate impact, *California Law Review* 104 (2016) 671.
- [6] J. Buolamwini, T. Gebru, Gender shades: Intersectional accuracy disparities in commercial gender classification, in: *FAT*, 2018.
- [7] S. U. Noble, *Algorithms of oppression: How search engines reinforce racism*, 2018.
- [8] R. Baeza-Yates, Bias on the web and beyond: an accessibility point of view, *Proceedings of the 17th International Web for All Conference* (2020).
- [9] R. Baeza-Yates, Biases on social media data: (keynote extended abstract), *Companion Proceedings of the Web Conference 2020* (2020).

- [10] E. Stamboliev, T. Christiaens, How empty is trustworthy ai? a discourse analysis of the ethics guidelines of trustworthy ai, *Critical Policy Studies* 0 (2024) 1–18. URL: <https://doi.org/10.1080/19460171.2024.2315431>. doi:10.1080/19460171.2024.2315431. arXiv:<https://doi.org/10.1080/19460171.2024.2315431>.
- [11] H. H. L. E. Group), Ethics guidelines for trustworthy ai., 2019. URL: <https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html>. doi:<https://ec.europa.eu/futurium/en/ai-alliance-consultation.1.html>.
- [12] J. M. Alvarez, A. B. Colmenarejo, A. Elobaid, S. Fabbri, M. Fahimi, A. Ferrara, S. Ghodsi, C. Mougán, I. Papageorgiou, P. Reyero, M. Russo, K. M. Scott, L. State, X. Zhao, S. Ruggieri, Policy advice and best practices on bias and fairness in ai, *Ethics and Information Technology* (2024). URL: <https://doi.org/10.1007/s10676-024-09746-w>. doi:10.1007/s10676-024-09746-w.
- [13] R. Schwartz, A. Vassilev, K. K. Greene, L. Perine, A. Burt, P. Hall, Towards a standard for identifying and managing bias in artificial intelligence, 2022. URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=934464. doi:<https://doi.org/10.6028/NIST.SP.1270>.
- [14] N. Mehrabi, F. Morstatter, N. A. Saxena, K. Lerman, A. G. Galstyan, A survey on bias and fairness in machine learning, *ACM Computing Surveys (CSUR)* 54 (2019) 1 – 35.
- [15] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. III, K. Crawford, Datasheets for datasets, *Commun. ACM* 64 (2021) 86–92. URL: <https://doi.org/10.1145/3458723>. doi:10.1145/3458723.
- [16] I. D. Raji, J. Yang, About ml: Annotation and benchmarking on understanding and transparency of machine learning lifecycles, *ArXiv abs/1912.06166* (2019).
- [17] J. Stoyanovich, S. Abiteboul, B. Howe, H. V. Jagadish, S. Schelter, Responsible data management, *Commun. ACM* 65 (2022) 64–74. URL: <https://doi.org/10.1145/3488717>. doi:10.1145/3488717.
- [18] I. D. Raji, A. Smart, R. N. White, M. Mitchell, T. Gebru, B. Hutchinson, J. Smith-Loud, D. Theron, P. Barnes, Closing the ai accountability gap: Defining an end-to-end framework for internal algorithmic auditing, in: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT* '20*, ACM, 2020, p. 33–44. URL: <https://doi.org/10.1145/3351095.3372873>. doi:10.1145/3351095.3372873.
- [19] E. M. Bender, B. Friedman, Data statements for natural language processing: Toward mitigating system bias and enabling better science, *Transactions of the Association for Computational Linguistics* 6 (2018) 587–604. URL: <https://aclanthology.org/Q18-1041>. doi:10.1162/tacl_a_00041.
- [20] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, T. Gebru, Model cards for model reporting, in: *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* '19*, ACM, 2019, p. 220–229. URL: <https://doi.org/10.1145/3287560.3287596>. doi:10.1145/3287560.3287596.
- [21] I. Fernandez, C. Aceta, E. Gilabert, I. Esnaola-Gonzalez, Fides: An ontology-based approach for making machine learning systems accountable, *Journal of Web Semantics* 79 (2023) 100808. URL: <https://www.sciencedirect.com/science/article/pii/S1570826823000379>. doi:<https://doi.org/10.1016/j.websem.2023.100808>.
- [22] P. Reyero-Lobo, E. Daga, H. Alani, M. Fernández, Semantic web technologies and bias in

artificial intelligence: A systematic literature review, 2022.

- [23] N. Noy, C. Goble, Are we cobblers without shoes? making computer science data fair, *Commun. ACM* 66 (2022) 36–38. URL: <https://doi.org/10.1145/3528574>. doi:10.1145/3528574.
- [24] W. Mark, et al., The fair guiding principles for scientific data management and stewardship, *Scientific data* 3 (2016) 1–9.
- [25] J. S. Franklin, K. Bhanot, M. Ghalwash, K. P. Bennett, J. McCusker, D. L. McGuinness, An ontology for fairness metrics, in: *Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, AIES '22*, Association for Computing Machinery, New York, NY, USA, 2022, p. 265–275. URL: <https://doi.org/10.1145/3514094.3534137>. doi:10.1145/3514094.3534137.
- [26] J. S. Franklin, H. Powers, J. S. Erickson, J. P. McCusker, D. L. McGuinness, K. P. Bennett, An ontology for reasoning about fairness in regression and machine learning, in: *Iberoamerican Conference on Knowledge Graphs and Semantic Web*, 2023.
- [27] D. Golpayegani, H. J. Pandit, D. Lewis, AIRO: an ontology for representing AI risks based on the proposed EU AI act and ISO risk management standards, in: A. Dimou, S. Neumaier, T. Pellegrini, S. Vahdati (Eds.), *Towards a Knowledge-Aware AI - SEMANTiCS 2022 - Proceedings of the 18th International Conference on Semantic Systems*, 13-15 September 2022, Vienna, Austria, volume 55 of *Studies on the Semantic Web*, IOS Press, 2022, pp. 51–65. URL: <https://doi.org/10.3233/SSW220008>. doi:10.3233/SSW220008.
- [28] D. Golpayegani, H. J. Pandit, D. Lewis, To be high-risk, or not to be—semantic specifications and implications of the ai act’s high-risk ai applications and harmonised standards, in: *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT '23*, Association for Computing Machinery, New York, NY, USA, 2023, p. 905–915. URL: <https://doi.org/10.1145/3593013.3594050>. doi:10.1145/3593013.3594050.
- [29] N. F. Noy, C. A. Goble, Are we cobblers without shoes?: Making computer science data FAIR, *Commun. ACM* 66 (2023) 36–38. URL: <https://doi.org/10.1145/3528574>. doi:10.1145/3528574.
- [30] T. R. Gruber, Toward principles for the design of ontologies used for knowledge sharing?, *International Journal of Human-Computer Studies* 43 (1995) 907–928. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1071581985710816>. doi:10.1006/ijhc.1995.1081.
- [31] A. M. Kaushik, R. Mutharaju, Chapter 21. an ontology design pattern for modeling bias, in: *Studies on the Semantic Web*, IOS Press, 2021. URL: <https://doi.org/10.3233/ssw210024>. doi:10.3233/ssw210024.
- [32] E. Pitoura, Social-minded measures of data quality: Fairness, diversity, and lack of bias, *J. Data and Information Quality* 12 (2020). URL: <https://doi.org/10.1145/3404193>. doi:10.1145/3404193.
- [33] N. Shahbazi, Y. Lin, A. Asudeh, H. V. Jagadish, Representation bias in data: A survey on identification and resolution techniques, *ACM Comput. Surv.* 55 (2023). URL: <https://doi.org/10.1145/3588433>. doi:10.1145/3588433.
- [34] A. Aler Tubella, D. Coelho Mollo, A. Dahlgren Lindström, H. Devinney, V. Dignum, P. Ericson, A. Jonsson, T. Kampik, T. Lenaerts, J. A. Mendez, J. C. Nieves, Acrocpolis: A descriptive framework for making sense of fairness, in: *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT '23*, ACM, 2023, p. 1014–1025. URL: <https://doi.org/10.1145/3593013.3594059>. doi:10.1145/3593013.3594059.

- [35] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV Ontology, W3C Recommendation, World Wide Web Consortium, United States, 2013.
- [36] G. C. Publio, D. Esteves, A. Ławrynowicz, P. Panov, L. Soldatova, T. Soru, J. Vanschoren, H. Zafar, ML-schema: Exposing the semantics of machine learning with schemas and ontologies, 2018. URL: <https://arxiv.org/abs/1807.05351>. doi:10.48550/ARXIV.1807.05351.
- [37] R. Albertoni, S. Colantonio, P. Skrzypczynski, J. Stefanowski, Reproducibility of machine learning: Terminology, recommendations and open issues, *ArXiv abs/2302.12691* (2023).
- [38] R. Albertoni, D. Browning, S. J. D. Cox, A. N. Gonzalez-Beltran, A. Perego, P. Winstanley, The w3c data catalog vocabulary, version 2: Rationale, design principles, and uptake, *ArXiv abs/2303.08883* (2023).
- [39] I. Hupont, D. Fernández-Llorca, S. Baldassarri, E. Gómez, Use case cards: a use case reporting framework inspired by the european ai act, *Ethics and Information Technology* 26 (2024). URL: <https://doi.org/10.1007/s10676-024-09757-7>. doi:10.1007/s10676-024-09757-7.
- [40] C. Sun, A. Asudeh, H. V. Jagadish, B. Howe, J. Stoyanovich, Mithralabel: Flexible dataset nutritional labels for responsible data science, in: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19, 2019*, p. 2893–2896. URL: <https://doi.org/10.1145/3357384.3357853>. doi:10.1145/3357384.3357853.
- [41] A. Wang, A. Liu, R. Zhang, A. Kleiman, L. Kim, D. Zhao, I. Shirai, A. Narayanan, O. Ruskovskiy, Revise: A tool for measuring and mitigating bias in visual datasets, *Int. J. Comput. Vision* 130 (2022) 1790–1810. URL: <https://doi.org/10.1007/s11263-022-01625-5>. doi:10.1007/s11263-022-01625-5.
- [42] G. P. . A. Research, Know your data, 2021. URL: <https://knowyourdata.withgoogle.com/docs/>, access:10.06.2022.
- [43] H. F. Research, Data Measurements Too, <https://huggingface.co/spaces/huggingface/data-measurements-tool>, 2022.
- [44] F. AI, Fairness flow, <https://ai.facebook.com/blog/how-were-using-fairness-flow-to-help-build-ai-that-works-better-for-everyone/>, 2021.
- [45] E. Pitoura, Social-minded measures of data quality, *Journal of Data and Information Quality* 12 (2020) 1–8. URL: <https://doi.org/10.1145/3404193>. doi:10.1145/3404193.
- [46] A. Balayn, M. Yurrita, J. Yang, U. Gadiraju, “fairness toolkits, a checkbox culture?” on the factors that fragment developer practices in handling algorithmic harms, in: *Proceedings of the 2023 AAI/ACM Conference on AI, Ethics, and Society, AIES '23, ACM, 2023*, p. 482–495. URL: <https://doi.org/10.1145/3600211.3604674>. doi:10.1145/3600211.3604674.
- [47] R. M. Shelby, S. Rismani, K. Henne, A. Moon, N. Rostamzadeh, P. Nicholas, N. F. Yilla, J. Gallegos, A. Smart, E. Garcia, G. Virk, Sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction, 2022.
- [48] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, R. S. Zemel, Fairness through awareness, in: *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012, ACM, 2012*, pp. 214–226. doi:10.1145/2090236.2090255.
- [49] P. Rezero Lobo, J. Kwarteng, M. Russo, M. Fahimi, K. Scott, A. Ferrara, I. Sen, M. Fernandez, A multidisciplinary lens of bias in hate speech, in: *Proceedings of the International Conference on Advances in Social Networks Analysis and Mining, ASONAM*

- '23, Association for Computing Machinery, New York, NY, USA, 2024, p. 121–125. URL: <https://doi.org/10.1145/3625007.3627491>. doi:10.1145/3625007.3627491.
- [50] M. Miceli, T. Yang, L. Naudts, M. Schuessler, D. Serbanescu, A. Hanna, Documenting computer vision datasets: An invitation to reflexive data practices, in: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21, ACM, 2021, p. 161–172. URL: <https://doi.org/10.1145/3442188.3445880>.
- [51] A. K. Heger, L. B. Marquis, M. Vorvoreanu, H. Wallach, J. Wortman Vaughan, Understanding machine learning practitioners' data documentation perceptions, needs, challenges, and desiderata, Proc. ACM Hum.-Comput. Interact. 6 (2022). URL: <https://doi.org/10.1145/3555760>. doi:10.1145/3555760.
- [52] S. Fabbrizzi, S. Papadopoulos, E. Ntoutsis, I. Kompatsiaris, A survey on bias in visual datasets, Comput. Vis. Image Underst. 223 (2022). URL: <https://doi.org/10.1016/j.cviu.2022.103552>. doi:10.1016/j.cviu.2022.103552.
- [53] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, A. Galstyan, A survey on bias and fairness in machine learning, ACM Comput. Surv. 54 (2021). URL: <https://doi.org/10.1145/3457607>. doi:10.1145/3457607.
- [54] A. Olteanu, C. Castillo, F. Diaz, E. Kiciman, Social data: Biases, methodological pitfalls, and ethical boundaries, Frontiers in Big Data 2 (2019). URL: <https://www.microsoft.com/en-us/research/publication/social-data-biases-methodological-pitfalls-and-ethical-boundaries/>.
- [55] E. F. Kendall, D. L. McGuinness, Ontology Engineering, Synthesis Lectures on the Semantic Web: Theory and Technology, Morgan & Claypool Publishers, 2019. URL: <https://doi.org/10.2200/S00834ED1V01Y201802WBE018>. doi:10.2200/S00834ED1V01Y201802WBE018.
- [56] A. J. Miles, S. Bechhofer, Skos simple knowledge organization system reference, 2009. URL: <https://api.semanticscholar.org/CorpusID:58835891>.
- [57] L. Yu, Foaf: Friend of a friend, 2011. URL: <https://api.semanticscholar.org/CorpusID:60893017>.
- [58] R. Albertoni, A. Isaac, Introducing the data quality vocabulary (dqv), Semantic Web 12 (2020) 81–97.
- [59] S. L. Blodgett, S. Barocas, H. Daumé III, H. Wallach, Language (technology) is power: A critical survey of “bias” in NLP, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Online, 2020, pp. 5454–5476. URL: <https://aclanthology.org/2020.acl-main.485>. doi:10.18653/v1/2020.acl-main.485.
- [60] I. D. Raji, I. E. Kumar, A. Horowitz, A. Selbst, The fallacy of ai functionality, in: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22, Association for Computing Machinery, New York, NY, USA, 2022, p. 959–972. URL: <https://doi.org/10.1145/3531146.3533158>. doi:10.1145/3531146.3533158.
- [61] M. Färber, F. Bartscherer, C. Menne, A. Rettinger, Linked data quality of dbpedia, freebase, opencyc, wikidata, and yago, Semantic Web 9 (2018) 77–129.
- [62] M. Poveda-Villalón, A. Gómez-Pérez, M. C. Suárez-Figueroa, OOPS! (Ontology Pitfall Scanner!): An On-line Tool for Ontology Evaluation, International Journal on Semantic Web and Information Systems (IJSWIS) 10 (2014) 7–34.