Can you trust Wikidata?

Veronica Santos [a,*], Daniel Schwabe [b,c] and Sérgio Lifschitz [a]
[a] *BioBD Lab, Informatics Department, Pontifical Catholic University of Rio de Janeiro, RJ, Brazil*
E-mails: vdsantos@inf.puc-rio.br, sergio@inf.puc-rio.br
[b] *Associate Researcher at BioBD Lab, Informatics Department, Pontifical Catholic University of Rio de Janeiro, RJ, Brazil*
[c] *Visiting Researcher at AI Lab, Joszef Stefan Institute, Slovenia*
E-mail: dschwabe@gmail.com

Abstract.
In order to use a value retrieved from a Knowledge Graph (KG) for some computation, the user should, in principle, ensure that s/he trusts the veracity of the claim, i.e., considers the statement as a fact. Crowdsourced KGs, or KGs constructed by integrating several different information sources of varying quality, must be used via a trust layer. The veracity of each claim in the underlying KG should be evaluated, considering what is relevant to carrying out some action that motivates the information seeking. The present work aims to assess how well Wikidata (WD) supports the trust decision process implied when using its data. WD provides several mechanisms that can support this trust decision, and our KG Profiling, based on WD claims and schema, elaborates an analysis of how multiple points of view, controversies, and potentially incomplete or incongruent content are presented and represented.

Keywords: Trust, Contextual, KG Profiling

1. Introduction - Data in Wikidata

Wikidata [1] (WD) is currently one of the most extensive publicly available Knowledge Graphs (KGs). Numerous other websites are regularly using it, services (search engines, personal assistants, libraries, and museums [2], [3]), applications (Daimler, Lufthansa, Novartis, data journalists, ...), and research projects [4]. It is regarded by many as part of the Semantic Web (e.g., [5]).

In most applications using KGs, one needs to obtain the value of some property of an entity (item in WD parlance) to make some computation. For example, finding the capital city of a country or province; obtaining a person's birth or death date; determining the author of some artifact; obtaining some physical property of a known substance; etc.

Consider the following example. Suppose a person wants to claim some inheritance from Ana Maria Imeni by being a descendant of Menotti Garibaldi, the son of the famous revolutionary Giuseppe Garibaldi. According to WD, there are two claims for who his mother was; see figure 1.

By design, WD contains statements based on claims about items, e.g., `<USA capital Washington>`. In principle, there is no guarantee of the truthfulness of claims.

We have argued elsewhere that crowdsourced KGs, or KGs constructed by integrating several different information sources of varying quality, must be used via a trust layer. It is where a decision is made about the veracity of each claim in the underlying KG relevant to carrying out some action being contemplated [6, 7]. Consequently, the data user should have their trust criteria to accept a valid claim and use it in an intended computation.

*Corresponding author. E-mail: vdsantos@inf.puc-rio.br.

Fig. 1. Menotti Garibaldi mothers.

This additional layer, with some generalizations, has already been argued about statements in WD and the Semantic Web, in general, [8]. It has also been recognized in the context of *polyvocal KGs* [9] that incorporate multiple *voices* (or points-of-view) about a subject.

The present work aims to assess how well WD supports the trust decision process implied when using its data.

WD provides several mechanisms that can support this trust decision, allowing (i) attributes of properties, called qualifiers; (ii) references to statements to support their veracity; and (iii) statements to be ranked as preferred, normal, or deprecated. It also provides properties to record actors' disagreements about a statement's veracity explicitly. We will detail these mechanisms further in Section 2.

The remaining sections detail the mechanisms present in WD to support trust decisions. They examine the explicit support for disagreements about the veracity of statements. They also look at incompleteness and how well constraint violations capture them. We investigate the support for detecting incongruences and what additional information it provides to resolve them in a trust decision. We also look closely at the support for provenance from a broader perspective. Finally, we present an analysis of the observations made on how well WD supports a trust decision based on the mechanisms mentioned earlier, discuss related work, and point to ongoing and future work.

2. Motivation - Incongruences in Wikidata

WD allows multiple statements with different values for the same property of the same item, i.e., expressing different perspectives about a subject even if they contradict one another. While this may be entirely appropriate in some cases (e.g., `spouse`, `head of government`, etc.), it is not so for others, such as inverse functional properties (e.g., `mother`, `date of birth`, etc.). Even though the trust process applies to statements that assert a single value for a property of an item, it is more crucial in situations where there are multiple statements about a property of an item with different values. A single value is needed for some computation. In such cases, beyond applying the trust process to each asserted value, the user is further confronted with the other decision of which value to choose.

According to the Cambridge Dictionary, a controversy is *a disagreement, often a public one, that involves different ideas or opinions about something*; Merriam-Webster's definition is *a discussion marked especially by the expression of opposing views*, and Oxford English Dictionary defines it as *disagreement, typically when prolonged, public, and heated*. From these definitions, one could interpret the occurrence of statements with multiple different values for the same item in a KG as indicators of potential controversies.

A more careful examination of such occurrences indicates the existence of three different types of situations - incompleteness, incongruences, and controversies. We discuss each one next.

The values for specific properties are sometimes time-dependent (e.g., `position held`, `capital`) or location-dependent (e.g., `boiling point`), or, more generally, context-dependent. Such dependencies are represented in Wikidata through the use of Qualifiers [10], which provide contextual information for statements.

The intended semantics for a property are further characterized by specifying constraints over these properties, which formulate restrictions on how these properties should be used within statements in WD [1].

Constraints are rules on properties that specify how properties should be used. There are two types of constraints to check properties and their qualifiers: `required qualifiers` (Q21510856) and `allowed qualifiers` (Q21510851). For example, the position held should have `start time`, defined via a `required qualifier constraint` for this property, and possibly `end time` qualifiers, specified via an `allowed qualifier constraint` for this property.

Another constraint of interest is `property scope` (Q53869507) that specifies how a property should be used: `as main value`, `as qualifier`, or `as reference`. However, constraints in WD work as hints and are not automatically enforced. Consequently, it turns out that many statements in WD that should have such contextual information in actuality do not. We identified that only 21% of the WD claims have qualifications. Due to its crowdsourced and distributed authoring model, incompletenesses often creep into the WD KG [11],[12].

In this situation, multiple values in a query result are not necessarily controversial. It can be the result of an improperly (incompletely) formulated query or a result of a shortcoming in the data present in the KG.

A second type of situation is when either no context applies (e.g., for pure inverse-functional properties such as `date of birth`) or when even including contextual information, a query may result in multiple different values for the same item.

The truth of the other may imply the truth of one of the statements (e.g., `<George Washington place of burial Mount Vernon>` and `<George Washington place of burial Washington's Tomb>`, since Washington's Tomb is located in Mount Vernon), and there is no contradiction in both being true at the same time. On the other hand, there are situations in which the two (or more) values for that property cannot be true in the same context (i.e., for the same values of qualifiers, if available). We refer to this situation as an incongruence.

Following the dictionary definitions of *controversy* given previously, one can rephrase them, for our purposes, as an *incongruence that also involves a discussion and disagreement among authors of the statements involved, possibly over an extended period*. For crowdsourced online resources such as WD and Wikipedia, a controversy is manifested through the *talk* pages associated with an entry, as well as by patterns in their edit history (see [13],[14], [15], [16]). Such analysis is outside the scope of this paper and is left for future work; in this respect, we will focus on incomplete and incongruent content in WD.

3. The trust process using WD

We have already pointed out that to use a value stated in a WD claim for some computation, the user should, in principle, ensure that s/he trusts the veracity of the claim, i.e., considers the statement a fact. This points to the need for an additional layer above the KG, which we call the Trust Layer, where trust policies and decision criteria are represented. The trust decision is made (for each user) based on the intended use (action) for the data retrieved from WD and additional information about the context in which the statement should be considered. In WD, this additional information comes in three forms in WDs data model [5]:

- Qualifiers, statements about statements, better characterize the relation being asserted. For example, «`France capital Versailles> <start time 1682>, <end time 1789`», which states the temporal context of the statement `<France capital Versailles>` as being valid from 1682 to 1789. Here `start time` and `end time` are examples of qualifiers. In general, qualifiers provide context for statements in WD [2]. Qualifiers themselves cannot be further qualified.
- WD statements have a rank, which can be `normal`, `preferred`, or `deprecated`. The latter two are meant to indicate a community consensus about the veracity or falsehood of the statement. Although seldom used, it is possible to provide a reason for some statements having a deprecated and a preferred rank through the respective qualifiers.

[1] https://www.wikidata.org/wiki/Help:Property_constraints_portal
[2] https://www.wikidata.org/wiki/Help:Property_constraints_portal/Qualifiers

– Statements can also have references (or sources). A reference can be a URL of some external source or a link to a WD item that supports the claim being made. References cannot be associated with qualifiers of a statement since the WD data model does not follow a multi-layer graph model [17].

To allow a more precise terminology and avoid ambiguity, in this paper, we refer to the role of a property in a (`main value`) statement as **Predicate**; to a pair (property value) associated to a statement as a **Qualification**, and to the role of the property in a qualification as **Qualifier**; and a property appearing in (property value) pair in a reference associated with a statement as **Referrer**.

In addition to these *structural* mechanisms, WD also provides ontological support recording controversial statements by explicitly marking them with the `disputed by` Qualifier (P1310). Using this property as a qualifier explicitly indicates that some source (a person, organization, but also a document, or legal instrument) disputes the veracity of the qualified statement.

Given the above information, the user must apply their trust policies, taking these information items into account to decide whether to accept or reject the claim's veracity. An important criterion often used in trust policies is provenance - e.g., who made a claim, or how the claim was established. Provenance may be given through references or some of the qualifiers (e.g., `stated in`). However, there are often statements without references (see section 4.5) or qualifiers. In such cases, the source of the statement must be attributed to the user who asserted it in WD.

Unfortunately, the id of the user who asserted a statement is not directly available as an item in WD itself, although it can be retrieved by accessing the underlying database (Wikibase). This significantly hampers the ability to build a trust layer on top of WD alone, as trust policies relying on the authorship of a statement (claim) would not be able to be evaluated.

This paper will not elaborate further on how such a trust layer could be defined and implemented. Rather, we analyze the current status of WD to support the trust process by properly recording multiple points of view, explicitly representing controversies, characterizing potentially incomplete or incongruent content, and providing provenance information.

4. WD support for the trust process

This section presents WD profiling data concerning incompleteness, incongruences, ranking, references, and provenance. This data was compiled using the DWD dump provided by the KG Center at ISI [3], reflecting the WD dump of June 2022 [4].

First, we present some general statistics about WD to give an overview of this KG. There are 559,038,971 claims using 9,653 properties as predicates, and 141,983,745 qualifications are associated with claims using 9,906 properties as qualifiers. There were 10,089 properties, and 90% of them have a `property scope constraint` assigned to them, distributed according to table 1 [5].

Table 1
Property Scope Constraint Distribution

Property Scope	Property Count
as main value	8663
as reference	6058
none	1099
as qualifier	745

[3] https://drive.google.com/drive/folders/1a6cUI1UEWRTNbvqtLAfJU0wEJ4ssTqdz?usp=share_link

[4] Note that this dump does not include instances of `scholarly article` (Q13442814) and `review article` (Q7318358) and their subclasses, since their frequency (over 50% of all statements) would distort the statistics. It also does not include statements annotated with `deprecated rank`.

[5] All data collected, statistics, scripts, and queries used for conducting this research can be found in https://github.com/versant2612/WD-Profiling

We observe that only 363 properties have a `required qualifier` constraint specification, and 761 have a `allowed qualifier constraint` specification, representing respectively 3,6% and 7,5% among 10,089 properties. Figures 2 and 3 present the distribution for both predicates and qualifiers for the top 50 most frequent ones. Those predicates account for 62,25% of all statements, and qualifiers account for 93,5% of all qualifications. Table 2 and 3 detail the top-10 of each one. It is essential to mention that: (1) According to the Property Talk page [6], there are more than 108 million claims using P31 property, but our dump version is from June 2022. Furthermore, the dump excluded almost 50% of statements, as noted in a previous footnote (2) Qualifier `of` (P642 property), while being among the most used for qualifications, is slated to become deprecated, see [7].

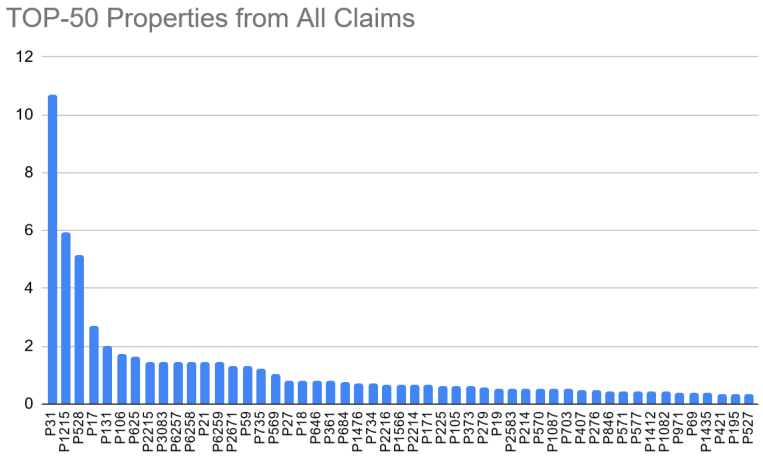

Fig. 2. Top-50 properties used as predicates from Claims.

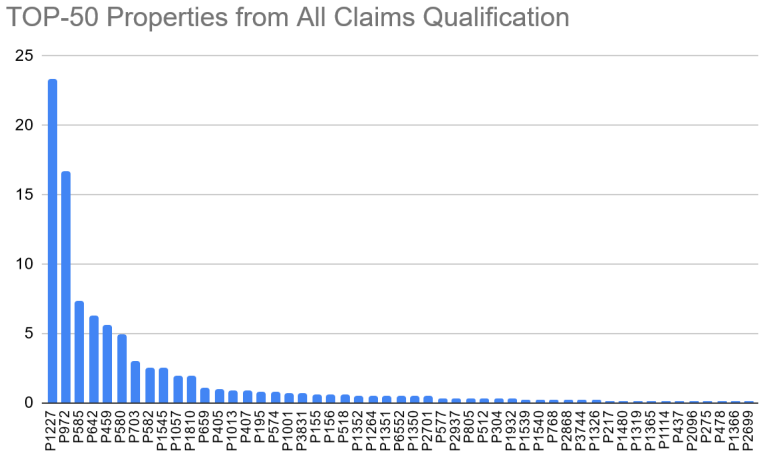

Fig. 3. Top-50 properties used as qualifiers from Qualifications.

[6] https://www.wikidata.org/wiki/Property_talk:P31 accessed in Feb 2023
[7] https://www.wikidata.org/wiki/Property:P642.

Table 2
Top-10 predicates in all claims. The frequency is relative to the total number of statements.

Predicate	Count	Label	Frequency	Accumulated Frequency
P31	59,717,980	instance of	10,68	10,68
P1215	33,122,376	apparent magnitude	5,92	16,6
P528	28,738,709	catalog code	5,14	21,74
P17	14,996,553	country	2,68	24,42
P131	11,371,144	located in the administrative territorial entity	2,03	26,45
P106	9,608,349	occupation	1,72	28,17
P625	9,267,000	coordinate location	1,66	29,83
P2215	8,207,685	proper motion	1,47	31,3
P3083	8,150,658	SIMBAD ID	1,46	32,76
P6257	8,091,255	right ascension	1,45	34,21

Table 3
Top-10 qualifiers in all qualifications. The frequency is relative to the total number of qualifications.

Qualifier	Count	Label	Frequency	Accumulated Frequency
P1227	33,122,324	astronomical filter	23,33	23,33
P972	23,776,643	catalog	16,75	40,08
P585	10,432,968	point in time	7,35	47,43
P642	8,966,228	of	6,31	53,74
P459	7,930,570	determination method	5,59	59,33
P580	7,048,511	start time	4,96	64,29
P703	4,317,830	found in taxon	3,04	67,33
P582	3,601,028	end time	2,54	69,87
P1545	3,597,770	series ordinal	2,53	72,4
P1057	2,776,317	chromosome	1,96	74,36

4.1. Explicit record of disagreements

We examine here the detailed record of disagreements via the use of the `disputed by qualifier` (P1310), which indicates explicitly that some source disputes the veracity of the qualified statement.

In an ideal disagreement record, each source should state what it claims is the correct value for the property being disputed. For example, `Jerusalem` (Q1218) is the `capital of` (P1376) `Israel` (Q801), this statement is ranked as preferred and disputed by `United Nations` (Q1065). Another statement asserts that `Jerusalem` (Q1218) is the `capital of` (P1376) `State of Palestine` (Q219060), this statement has normal rank, and `applies to part` (P518) `East Jerusalem` (Q212938) and is `disputed by` `United States of America` (Q30) and `Israel` (Q801), see figure 4.

Interestingly, however, there are many cases in which a disagreement is recorded using this qualifier, but no alternative value is claimed. For example, the `Republic of Crimea` (Q15966495) is claimed to be an `instance of` (P31) `republic of Russia` (Q41162) - but no indication as to who is making this claim - which is `disputed by` (P1310) both the `Parliamentary Assembly of the Council of Europe` (Q939743) and the `United Nations General Assembly` (Q47423). There are no claims by the latter about what `Republic of Crimea` is an `instance of`, distinct from `republic of Russia`, see figure 5.

There is another way to represent a `disputed by` assigning a `sourcing circumstance` (P1480) qualifier to a statement with a `value disputed` (Q18912752). We ignore this usage since their frequency is shallow.

The `disputed by` property is also used as a predicate to directly represent an opposition of concepts rather than a disagreement between statements. For example, the opposition of `Capitalism` (Q6206) and `Marxism`

Fig. 4. An example of a dispute with multiple points of views.

Fig. 5. An example of an one-sided disagreement.

(Q7264). We will disregard this type of disagreement, as it is different from the phenomena we are interested in. Furthermore, this use violates the `scope constraint` on the `disputed by` property.

`Disputed by` *Frequencies*

Next, we present statistics about explicit disputes in WD. There was 1577 [8] claims out of 559,038,971 in total, having `disputed by` as qualifier, representing 0,0003%. We also found that only 140 different properties, out of 9653 in total, were used as the predicate in these claim statements. In figure 6, we show the top-50 most frequent properties, corresponding to 92% of statements with `disputed by` qualifiers. Nearly 77% among the ten most frequent are properties about territory, medicine, diseases, people, and fiction (see table 4).

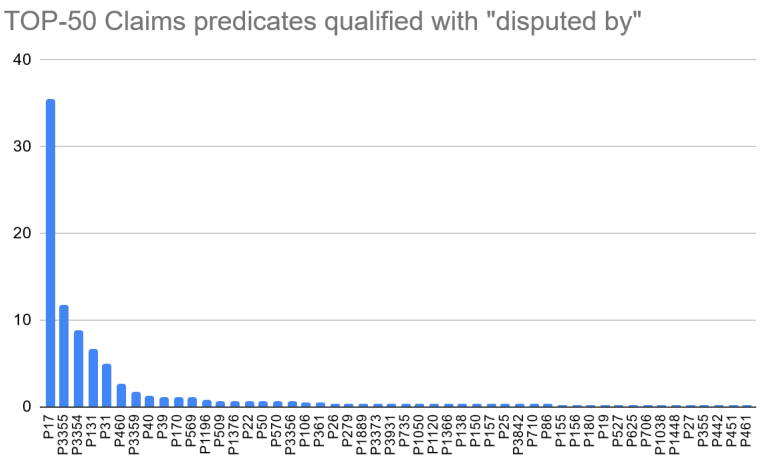

Fig. 6. Top-50 most frequent properties in Claims qualified with `disputed by`.

Table 4
Property types for the Top-10 Claims predicates qualified with `disputed by`

Property Label	Property Type
child	Wikidata property for human relationships
creator	Wikidata property associated with people Wikidata property for items about fictional characters Wikidata property related to creative works
position held	Wikidata property for items about people
instance of	Wikidata property for the relationship of the element to its class
positive therapeutic predictor for	Wikidata property related to disease biomarkers Wikidata property related to medicine
negative therapeutic predictor for	Wikidata property related to disease biomarkers Wikidata property related to medicine
negative prognostic predictor for	Wikidata property related to disease biomarkers Wikidata property related to medicine
country	Wikidata property related to places
said to be the same as	Wikidata property to describe the elements of identity
located in the administrative territorial entity	Wikidata property to indicate a location

This profiling information indicates that only a tiny fraction of statements in WD record explicit disagreements (0,0003%), perhaps unsurprisingly involving topics that are already controversial in the physical world. The statements capturing these disagreements often do not use the available properties in a way that can support a user in

[8] According to the Property Talk page https://www.wikidata.org/wiki/Property_talk:P1310 accessed in Feb. 2023, more than 2246 qualifications are using the P1310 property.

deciding which statement to accept as a fact if they must do so. This is further hampered by limitations of the WD data model, for example, needing to give some provenance information (or reference) to a qualifier statement (since qualifiers cannot be qualified). For instance, it might be essential to have a reference to the qualifier statement `supported by United States of America` used for the statement `<Jerusalem the capital of Israel>` (see Fig.4). Another hampering factor is the lack of author information for a statement as data in WD itself.

4.2. *Incompleteness of data/constraint violations*

Qualifiers and constraints are essential mechanisms to provide a context in WD. Looking at all statements, there are 141,983,745 qualifications over 116,211,413 claims, using 5,082 properties as predicates; qualifications were associated with claims using 9,905 properties as qualifiers. This means that only roughly 21% of statements provide some context information.

One can identify contextually incomplete statements by looking at violations of the required qualifier constraint. A `required qualifier constraint` can have a `constraint status`, which in turn can be `mandatory` or `suggested`. For simplicity, we refer to those `required constraints` with `constraint status mandatory` as `mandatory required qualifier constraint`, and likewise for `suggested required qualifier constraint`. The statistics presented next characterize them.

We found 359,445 statements with a `mandatory required qualifier` violation, representing 0,0643% of all claims. In comparison, 3,448,673 statements violated the `suggested required qualifier` constraint of their predicates, meaning nearly 0,62% of all claims. Figures 7 and 8 show the top-10 most frequent predicates used in those claims.

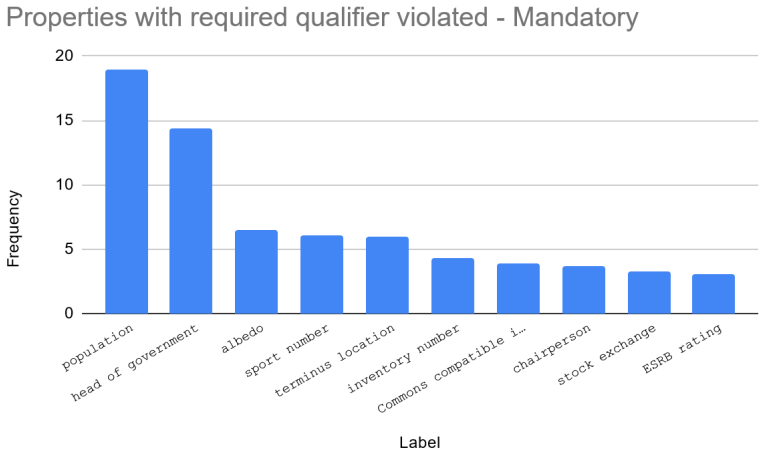

Fig. 7. Top-10 `Required Qualifier Constraint` Violations: Mandatory.

There were 395,604 `mandatory required qualifier` absences and 3,478,979 `suggested required qualifier` absences. Tables 5 and 6 show the top-10 missing qualifiers of each set.

The qualifiers `Point in time`, `Start Time` (related to temporal context), `Content Description`, and `Determination Method` are present in both sets. Indeed, the temporal qualifiers make up 51% of the mandatory ones and 22% of the suggested ones.

Some properties, defined as having a `restrictive qualifier` (Q61719275) type, must be used as qualifiers to restrict or modify the referent. Their absence may turn the statement inaccurate or meaningless. Among the missing `required qualifiers`, we found eight that correspond to restrictive qualifiers, and table 7 shows their number of violations. It should be noticed that the qualifier `applies to jurisdiction` (P1001) is an illustrative example of location-related `restrictive qualifier` since it restricts the statement in terms of territorial jurisdiction: a country, state, municipality, and so on.

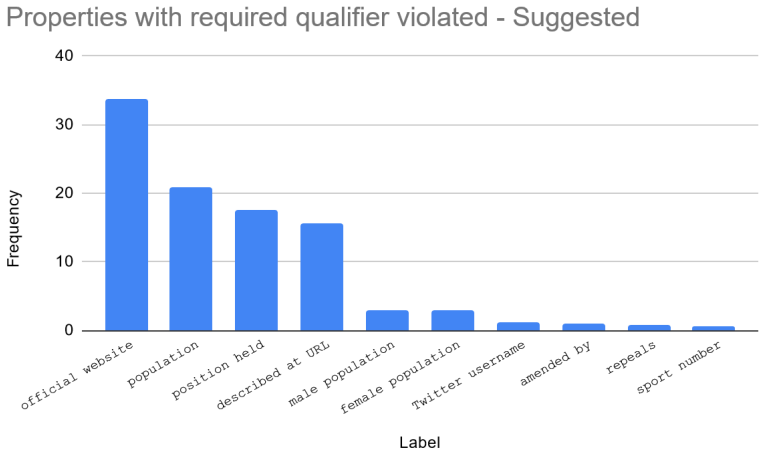

Fig. 8. Top-10 `Required Qualifier Constraint` Violations: Suggested.

Table 5

Top-10 Missing Qualifiers mentioned in required qualifier mandatory constraint. The frequency is relative to the total number of Qualifications absence. Highlighted predicates involve time.

Qualifier	Label	Count	Frequency	Accumulated Frequency
P585	`point in time`	112128	28,34	28,34
P580	`start time`	50230	12,7	41,04
P3415	`start period`	38268	9,67	50,71
P1013	criterion used	24668	6,24	56,95
P642	of	22462	5,68	62,63
P560	direction	21502	5,44	68,07
P195	collection	15414	3,9	71,97
P7367	content descriptor	12795	3,23	75,2
P2093	author name string	10896	2,75	77,95
P459	determination method	9929	2,51	80,46

4.3. Ranking

Sometimes an item should have only one value related to a property, which ideally is expressed through the `single-best-value constraint` [9]. Still, in practice, there are occurrences of multiple values. The rank of the statements, together with a qualifier, can be used to define which statement is the "best-one" and why. In the analyzed dataset, there were 5,480,866 preferred ranked claims, nearly 1% of all claims, and only 71,609 have the qualifier `reason for preferred` (P7452) [10].

Table 8 presents the top-10 most frequent properties in preferred ranked claims without the reason qualifier. We observe that demography properties are time-dependent and, for this reason, should not be specified with a `single-value constraint`, but rather the `single-best-value` seems to be more appropriate. In cases when the `point in time`, qualifier about these statements are unknown, the most recent one can be annotated as `preferred`. When the temporal context of claims is known, the annotation of `preferred` reinforces the notion of current or at least more recent value, as shown in figure 9 about Brazil's population in the year 2021.

[9] https://www.wikidata.org/wiki/Help:Property_constraints_portal/Single_best_value

[10] According to the Property Talk page https://www.wikidata.org/wiki/Property_talk:P7452 accessed in February 2023, there are approximately 122.000 qualifications.

Table 6

Top-10 Missing Qualifiers mentioned in required qualifier suggested constraint. The frequency is relative to the total number of Qualifications absence. Highlighted predicates involve time.

Qualifier	Label	Count	Frequency	Accumulated Frequency
P407	language of work or name	1717233	49,36	49,36
P459	determination method	926034	26,62	75,98
P580	`start time`	656267	18,86	94,84
P585	`point in time`	103465	2,97	97,81
P6552	Twitter user numeric ID	19426	0,56	98,37
P2676	rating certificate ID	13329	0,38	98,75
P972	catalog	8662	0,25	99
P137	operator	5202	0,15	99,15
P7367	content descriptor	4085	0,12	99,27
P4390	mapping relation type	3561	0,1	99,37

Table 7

Missing Restrictive Qualifiers from Mandatory Required Qualifiers Constraints

Qualifier	Label	Count (All Qualifications)	Count (Multi Value Qualification)
P3415	start period	38268	17555
P1013	criterion used	24668	16
P642	of	22462	592
P2210	relative to	91	4
P1552	has quality	39	0
P1001	`applies to jurisdiction`	13	2
P518	applies to part	8	5
P1264	valid in period	4	0

Ranking statements, as preferred, can be regarded as a mechanism to indicate a community consensus about the veracity (or at least the likelihood) of a statement. From the trust process point of view, it is a proxy for a more detailed determination process that each user goes through and does not consider for what purpose that claim will be used. This is supported by observing, for example, that nearly 72% of ranked statements (for `apparent magnitude`) are associated with a specific community, in this case, astronomy.

The statistics show that this mechanism is seldom used (1% of the statements) and even less frequently justified. Consequently, if a user must choose among multiple possible values for a property, s/he would rarely be able to rely on ranking. Furthermore, little support is given to help them determine if the criteria used by the community match their own for the intended action.

4.4. Incongruences

Incongruences among statements where multiple values are present for the same WD item and property should be detected after evaluating their values and contextual information. For instance, if someone is interested in who is the `head of state` of `Brazil` (Q155), a query looking only at the statements (ignoring qualifiers) will retrieve the last eight presidents in WD. However, all statements have the `required qualifier` `start time` (P580) as well as the `allowed qualifiers` `position held` (P39) and `determination method` (P459). Seven statements have the `allowed qualifier` `end time` (P582). By taking into account the contextual information (in this case, `start time` and `end time`, it is clear that there is no incongruence between the statements, even though they all refer to the `head of state` property of Brazil.

Nevertheless, there are many instances of statements that assert multiple different values for the same property of the same item and do not provide any qualifier to differentiate among them. Such statements indicate either a gap in the recorded information in WD or potential incongruences.

Table 8
Property types and frequency for the Top-10 preferred ranked claims The frequency is relative to the total number of Preferred Claims.

Property Label	Frequency %	Constraint	Property Type
`apparent magnitude`	72,37	single-value	Wikidata property for items about astronomical objects
population	4,17	single-value	Wikidata property for number of people, Wikidata property related to demography
social media followers	3,23		Wikidata property for number of people, Wikidata property related to online communities
located in the administrative territorial entity	3,15		Wikidata property to indicate a location
country	1,63		Wikidata property for items about places
instance of	1,43		Wikidata property for the relationship
contains administrative territorial entity	1,14		Wikidata property for items about places
OKTMO ID	1,1	single-value	Wikidata property for authority control for administrative subdivisions
male population	0,94	single-value	Wikidata property implying third-parties' gender, Wikidata property for number of people, Wikidata property related to demography

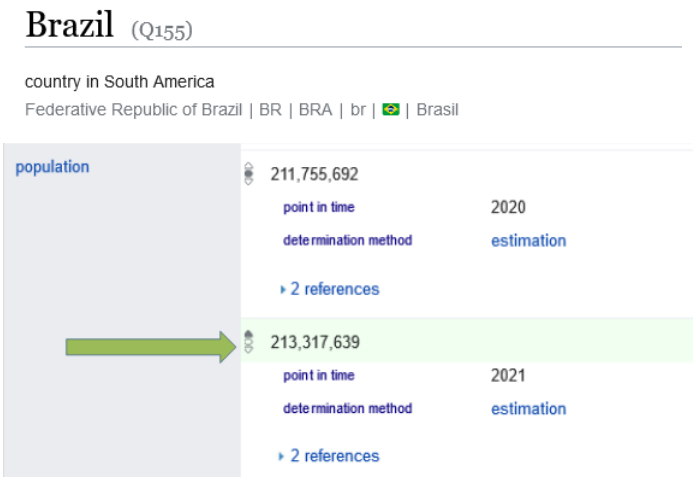

Fig. 9. Brazil's population of 2020 and 2021 where 2021 is marked as preferred (green).

Multiple Value Statements Statistics

Next, we present statistics about multiple value statements in WD and property constraints:

There are 132,552,453 statements with multiple values for the same item and the same property corresponding to 23,71% of all statements. We separate this set into statements with and without qualifiers. Of these, 91,148,503 have at least one qualification (approximately 16.3% of all statements). Overall, 1,333,012 statements with multiple values, representing 0.24% of all statements, violate one or more required qualifier constraints. Among the top-10 qualifier absences there are 3 time-related properties: `point in time` (P585): 36,74%, `start period` (P3415): 18,57% and `start time` (P580): 6,41%.

In a best-case scenario, where one assumes that all violations have been corrected, only 16.54% of all statements with multiple values would provide some contextual information. There are still 41,403,950 statements; 31,3% of those with multiple values (7,41% of all statements) do not have any qualifications.

We computed `required qualifier constraint` compliance based on the statement predicates, shown in Table 9. The top-10 qualifiers, present in table 9 correspond to nearly 19,2% of the total number of multiple

valued statements, with about 8% of them, involving qualifiers about time (`point in time`, `start time`).

Table 9

Top-10 `required qualifiers` in multiple valued statements. The frequency is relative to the total number of all statements with multiple values. Highlighted predicates involve time.

Qualifier	Label	Count	Frequency	Accumulated Frequency
P642	of	8209469	9,01	9,01
P585	`point in time`	6548388	7,18	16,19
P459	determination method	1524643	1,67	17,86
P580	`start time`	696591	0,76	18,62
P407	language of work or name	134829	0,15	18,77
P1352	ranking	131878	0,14	18,91
P1111	votes received	81761	0,09	19
P447	review score by	57777	0,06	19,06
P3294	encoding	44798	0,05	19,11
P195	collection	43065	0,05	19,16

We also checked the `allowed qualifier constraints` for the predicates in multiple valued statements. Once more, four qualifiers about time were among the top-10 most frequent, but here they correspond to only 26,5% of all multiple value statements. See table 10.

Table 10

Top-10 `allowed qualifiers` in multiple valued statements. The frequency is relative to the total number of all statements with multiple values. Highlighted predicates involve time.

Qualifier	Label	Count	Frequency	Accumulated Frequency
P972	catalog	22175343	24,33	24,33
P585	`point in time`	6715923	7,37	31,7
P580	`start time`	3101279	3,4	35,1
P582	`end time`	2358699	2,59	37,69
P459	determination method	1837514	2,02	39,71
P1545	series ordinal	1156594	1,27	40,98
P1264	`valid in period`	751061	0,82	41,8
P1351	number of points/goals/set scored	736338	0,81	42,61
P1350	number of matches played/races/starts	692587	0,76	43,37
P1013	criterion used	654304	0,72	44,09

It is important to note that not all cases of multiple values are inconsistent, as there may be a functional dependency or logical implication between them. Even though the statements are fully qualified and referred to, the incongruence may be only detected while applying the trust process to each asserted value according to task-specific criteria.

4.5. Provenance and References

Provenance in WD is mainly given through references. A reference is any property whose type is `Wikidata property to indicate a source` (Q18608359). This value can be an item in WD itself (e.g., a person or organization) or refer to an external source via URL. Some of these referrers (`reference properties`) can be used both as the `main value` of an item and as a qualifier in the role of reference to a statement. We focus only on their use as reference qualifiers in qualifications since, as such, they serve to provide provenance information.

Among the top-10 properties indicating the use of a source as qualifiers, shown in table 11 and representing 2,11% of all qualifications, seven have all three possible scopes (`as main value`, `as qualifier`, `as reference`). Strangely, `place of publication` (P291) does not have any `scope constraint`

at all; described by source (P1343) has scope as main value, and as qualifier but not as reference; and retrieved (P813) has scope as qualifier and as reference, but not as main value.

Table 11

Top-10 properties to indicate a source used as qualifiers. Frequency is calculated based on Qualifications using a property of type Wikidata property to indicate a source.

Qualifier	Count	Label	Frequency	Accumulated Frequency
P407	1242876	language of work or name	41,17	41,17
P577	537468	publication date	17,8	58,97
P304	441380	page(s)	14,62	73,59
P478	187030	volume	6,2	79,79
P291	105546	place of publication	3,5	83,29
P2093	98772	author name string	3,27	86,56
P1476	92759	title	3,07	89,63
P813	68490	retrieved	2,27	91,9
P1343	50777	described by source	1,68	93,58
P958	45364	section, verse, paragraph, or clause	1,5	95,08

We also observe that some properties are actually used as a referrer in references but aren't of the proper type (Q18608359). One often occurring example of this is main regulatory text (P92), which in addition only has as main value in a scope constraint (so it should not be used in a reference). Figure 10 about Republic of Crimea (Q15966495) shows an example of this misuse.

Fig. 10. Main Regulatory Text constraint scope violation shown in WD browser for the Republic of Crimea (Q15966495) item.

Because the ISI dump does not include references, we extracted them using the WD Virtuoso SPARQL Endpoint[11] that contains 1,428,322,719 statements and 103,740,680 references. We retrieved 241,831,575 referrer/value pairs that make up references considering only data types of values wikibase-item, url, external-id, and string since those represent provenance information. One reference can have several associated pairs, although, in practice, most references are composed of at most three pairs. We computed the frequency of each of the 5049 Referrers. We observed that seven of the top-10 most used referrers, shown in Table 12, are of type Wikidata property to indicate a source and the remaining three are external identifiers. Among the

[11] https://wikidata.demo.openlinksw.com/sparql accessed on Feb-04-2023

5049 properties used as referrers, there are 4105 whose data type is an `external identifier` corresponding to 17,9% as presented in Figure 11.

Table 12
Top-10 Properties used as Referrers in References. External Identifiers are highlighted.

Referrer	Count	Label	Types
P248	87044656	stated in	Wikidata property to indicate a source
P854	66044529	reference URL	Wikidata property to indicate a source
P698	30007333	`PubMed ID`	Wikidata property for authority control for works Wikidata property related to medicine Wikidata property for items about scholarly articles
P887	7312604	based on heuristic	Wikidata property to indicate a source non-restrictive qualifier
P4656	5658353	Wikimedia import URL	Wikidata property to indicate a source Wikidata property to link to Commons Wikidata property about Wikimedia entities
P143	5183797	imported from Wikimedia project	Wikidata property to indicate a source
P932	5123097	`PMCID`	Wikidata property for authority control for works Wikidata property for items about scholarly articles
P214	3846213	`VIAF ID`	Wikidata property for authority control for people Wikidata property for authority control for places Wikidata property for authority control for works Wikidata property to identify organizations Wikidata property for authority control, with reciprocal use of Wikidata Wikidata property for authority control by VIAF member 'Wikidata property for an identifier that does not imply notability Wikidata property widely reused by third-party entities Wikidata property for an identifier that generally has just one value, but rarely can have two or more values for the same entity
P1810	2525708	subject named as	Wikidata qualifier Wikidata property to indicate a name Wikidata property with datatype string that is not an external identifier property constraint
P351	2321925	`Entrez Gene ID`	Wikidata property related to biology Wikidata property to identify genes

It should be noted that provenance, in its broader sense, includes not only the sources but also the activities and other elements that are involved in the creation of an artifact (see, for example, the Prov-DM data model [12]). This means that other WD Properties can serve the role of provenance when used as qualifiers, e.g., `determination method` (P459), which has 7,930,570 occurrences as a qualifier. Still, this use is practically non-existent as a predicate (329 occurrences) and infrequently as a referrer (1,535,639 occurrences). As additional examples, the top 50 most frequent qualifiers include taxonomy-related properties defined as qualifiers using `property scope constraint`. For instance, `taxon author` (P405) - the author(s) that (optionally) may be cited with the scientific name, `year of taxon publication` (P574) - the year when this taxon was formally described or established, and `found in taxon` (P703) - the taxon in which the item can be found.

[12] https://www.w3.org/TR/2013/REC-prov-dm-20130430/

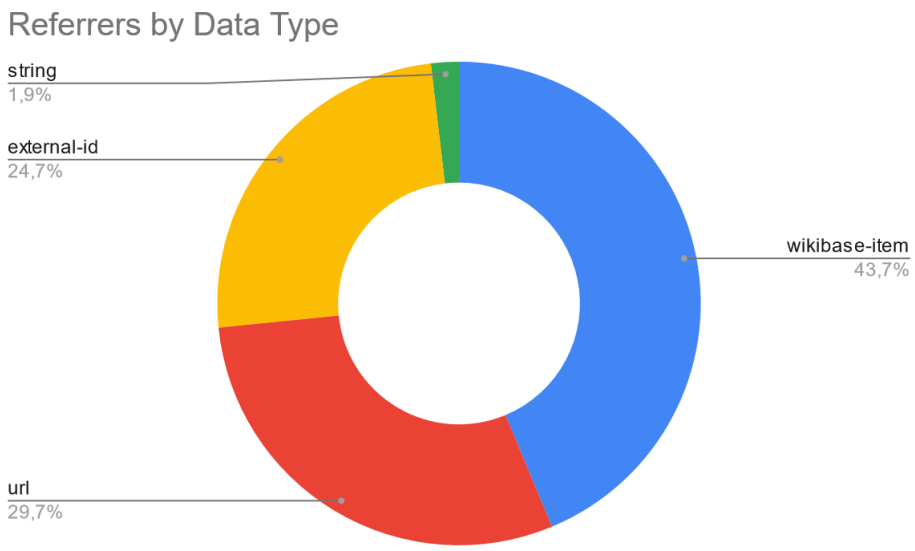

Fig. 11. Data type distribution of Referrers.

5. Conclusions

5.1. Discussion

To gain some insight into what these statistics tell us, there are several factors to consider, to whit:

- Some statements are missing qualifications because of a simple lack of information at creation time, and authors expect that either themselves or other users will eventually fill in this information when it becomes available, similar to Wikipedia missing links in page contents;
- Constraints for some properties are possibly mis-formulated - i.e., their intended semantics need to be understood by the user community that creates statements. In some sense, constraints reflect a community consensus about the semantics of properties that all users may not share. Furthermore, constraints do admit violations in some exceptional cases;
- Some properties require qualification, but this has yet to be captured in constraints when they were defined, despite the editorial process. This can also happen if the understanding of the semantics of the property evolves due to new knowledge being brought in or created;
- Finally, users are simply ignorant that certain properties should be qualified (although the editing interface warns them of this, and they are flagged in the standard WD browser;
- Incorrect or incomplete import scripts to automate the ingestion of new sources or unavailability of that information in the imported source.

We summarize here our observations about the various mechanisms provided by WD to support a trust decision about one of its claims.

Although the explicit recording of disagreement may seem like a natural way to capture these real-world situations in WD, only a negligible fraction of statements and a very small number of properties are present using this mechanism. The properties used in these statements involve mostly territory, medicine, diseases, and relations among people. The analysis also highlights a shortcoming of the WD data model, which prevents a qualifier from being further qualified or having a reference. A trust decision, therefore, cannot verify, for example, if some provenance information supports a qualification.

Considering that qualifications provide context information, we observe that only 21% of statements are qualified. Regarding `mandatory required qualifier constraint` violations (3,22% of all statements), one can see that even if they were to be satisfied, there would still be only 24,2% qualified statements.

`Suggested required qualifier constraint` violations have complex semantics to interpret since they are merely suggestions. Nevertheless, even if one considers the best case in which the data was provided to satisfy all of them (0,65% of all statements), there would still be only 24,8% qualified statements altogether.

This means that approximately 75,2% of statements do not have any context information via qualifiers. For those, any support for a trust decision would come from provenance information via authorship, which we already indicated is not provided in WD or through references.

We could not determine the total number of referenced statements in WD for computational reasons, as we had to extract those from the full dump. However, we can make some approximations based on the proportion of statements excluded from the ISI dump we used for all the other statistics. For this dump, it left out approximately 50% of all statements by ignoring scholarly and review articles and their sub-classes. Applying this proportion to the total reported in section 4.5, we estimate there are 51,870,340 references, of which only 43,7% refer to WD items, amounting to 22,667,338 references. As a lower bound, if we assume all of those references hypothetically apply to only those statements without qualification (which we know is not the case), this would still represent only 5,4% of all such statements. Therefore it is safe to conclude that at least 69,6% of all WD statements do not have any support for trust decisions.

As already pointed out throughout the paper, author information about statements is unavailable as part of the KG in WD. This would be the only possible information for those statements without qualifiers or references that can provide information to the trust process and therefore has a significant negative impact on WD's support for trust decisions.

Statements with multiple values for the same property for the same item can represent potential incongruences. For such statements, the trust decision must, in addition, filter each value claimed and often chose one of them to perform a computation. Again in a best-case scenario in which all constraint violations were to be corrected by providing missing values, there would still be 41,403,950 statements, 31,3% of those with multiple values (7,41% of all statements) without any qualifications.

It is worth noting that temporal information, often essential to determine the context of the veracity of a claim properly, is missing in 57% of claims that violate a time-related `mandatory required qualifier constraint`.

WD provides no, or at best very little, support to detect implicit disagreements, differentiate them from incomplete data, and support trust decisions.

In situations with multiple values, the user has to decide which value to choose beyond applying the trust process to each asserted value. WD provides the concept of ranking to aid the user in making this decision. We identified a shallow occurrence of ranked statements and even fewer with a qualifier to specify a `reason for preferred`, making it difficult to use the community consensus about the claims' veracity. Furthermore, multiple values are typical and expected for time-series types of information, and time-related qualifiers are crucial to dispel apparent incongruences. However, fully 61,7% of constraint violations of multiple-valued statements involve precisely time-related qualifiers.

The analysis shows that provenance is also available through qualifiers, and other types of properties related to methods and processes should also be considered as provenance. Even though 43,7% of references point to WD Items, there is still a large number of reference information pointing to external data sources using identifiers (24,7%), URLs (29,7%), and strings (1,9%). These external references typically require additional (often non-trivial) processing to be used in a trust decision about the claim where it is used as a reference. Furthermore, the WD model does not allow references to be associated with qualifiers since it is not a multi-layered model [18].

Looking at tools provided by Wikidata to support the trust process, one can look at ways to retrieve statements with preferred rank in case of multiple value statements. One such way proposed in WD is to use, when querying WD, statements annotated with preferred rank to retrieve the best non-deprecated ranked statement for a given property using the prefix `wdt`[13]. However, this applies only to predicates but not to qualifications, including the

[13] https://www.mediawiki.org/wiki/Wikibase/Indexing/RDF_Dump_Format

`reason for preferred` (P7452) qualifier, and references need reification. Thus, the information returned is incomplete in providing available information for the trust decision.

Another tool is that constraint violations can be visually verified using the standard WD browser interface, but it is very complex to infer this information at query time. A SPARQL query example for required qualifier violations, extracted from [19], is illustrated in [14].

From these analyses, it is possible that WD's support for trust decisions about its statements is low and could be improved significantly.

5.2. Related Work

Piscopo and Simperl (2019) [12] analyzed 28 previous papers about WD quality metrics using different dimensions and, similarly to Bizer et al. [8], they also considered that (i) data quality depends on the task and (ii) its relevance to a task is user-subjective. Although the authors stated that completeness is the dimension covered by the most significant number of papers, different from those, we chose to focus on qualification completeness based on required qualifiers violations. Such constraint violations may lead to incongruence and controversies in interpreting the data since it affects claims contextualization.

The authors of [11] developed a framework to detect and analyze low-quality statements in WD. They propose three metrics to identify WD Quality: Community-based indicators, Deprecation-based indicators, and Constraints-based indicators. Differently from those concerned with WD data quality in general, our analysis aims to detect WD incongruences, incompleteness, and controversy that affect criteria often used in trust policies, such as provenance and timeliness.

Regarding Provenance, [20] developed an approach to evaluate the relevance and authoritativeness of Wikidata references based on two complementary methods: microtask crowdsourcing and machine learning. Our analysis is complementary, looking at constraint violations and analysis of types of references.

WD Property Talk pages have useful information such as the property's current usage, query examples, and report about constraint violations. WD also has definitions of schemas in the form of shapes (subgraph patterns to describe a concept) using ShEx shape expressions. Some schemas, besides predicates with `p:` or `wdt:` prefixes, also uses qualifiers with `pq:` prefix (see `p:710` and its allowed qualifiers in `EntitySchema:E84` [15]) and referrers with `pr:` prefix (see `EntitySchema:E43` [16] and the `pr:P854` specification) to define a well-formed concept.

In [19], the authors proposed a set of SPARQL queries, formalizing all current WD property constraint specifications. These queries can be used to verify if a statement violates a constraint in online mode. Our profiling approach was developed using the Kypher language in the KGTK toolkit [21]. Still, we used the same query logic to verify some constraints violation of interest among all WD claims in batch mode. We complement our data using SPARQL queries to extract references data from an EndPoint.

Statistical methods to automatically build knowledge graphs create a simplified representation of knowledge where diverse perspectives tend to be resolved as conflicts. This approach introduces bias and loses traceability [9]. According to them, data is created based on a particular perspective or view. Data models that advocate being perspective-aware should be able to trace provenance in a broader sense, including creation and transformations.

WD, as a crowdsourcing effort where a voice corresponds to the data contributor perspective, enables the representation of provenance of claims at the individual level. Still, references are associated only with the data provider. The viewpoint of data contributors is unknown, and rank annotation of some statements reflects the consensus from a particular domain community. Complementary to this perspective, another viewpoint must be considered when building a contextualized knowledge graph to support the trusty layer: the data consumer. The Trust layer will apply trust policies to this perspective over data representing different viewpoints.

Bizer et al. [8] argued that web information should have its quality assessed according to task-specific criteria before being used. Their work is supported by Quality-based information filtering policies, that is, policies that users may choose for deciding if some piece of web information may be accepted or rejected to accomplish a specific

[14] https://w.wiki/6JNt

[15] https://www.wikidata.org/wiki/EntitySchema:E84

[16] https://www.wikidata.org/wiki/EntitySchema:E43

task. They developed the WIQA—Information Quality Assessment Framework composed of a set of software components dealing with uncertain information. Their proposed *quality layer* is similar to what we envisage for the trust layer for KGs to support the trust decision process implied when using its data.

Generating statistics on specific domains may have different characteristics from those we use in our approach. In [22], in addition to their tool to visualize research profile, the authors also present how WD has been used to store and disseminate bibliometric information and present scientometric statistics extracted from this information. Our approach can be applied to any sub-domain present in WD, such as Academics, Astronomy, Genomics, Biology in general, etc. Since the goal is to analyze WD support for the trust process through multiple points of view, controversies, and potentially incomplete or incongruent content.

5.3. Future work and Conclusions

Any open KG constructed collaboratively by contributors with different levels of domain knowledge and world views can suffer from wrong, biased, outdated, incomplete, and inconsistent content. The knowledge retrieved from a KG to be used to accomplish a specific task has to be more contextualized as possible to support the trust process. Here we adopted the context definition from [18]: *By context, we refer to the scope of truth, and thus talk about the context in which some data are held to be true*.

After analyzing the current status of WD, we realized that given the characteristics of the data model (such as the lack of references for qualifiers, incorrect specification of property scope in addition to the absence and ambiguities of definitions in constraints) and in the data (violations of constraints that generate incompleteness and possible interpretation errors) it is necessary to have an additional and separate layer, that will be user and task-dependent. And, to provide all information required by this layer, the knowledge retrieved from the KG should be explicitly contextualized, at least in terms of Provenance, Temporal, and Location dimensions.

As part of ongoing and future work, we are looking at ways to improve the ontology and the data model of a KG to support a trust layer better. This process will necessarily involve human participation through knowledge engineers, aided by automated and semi-automated tools. Even though this paper has explicitly focused on WD, we are looking at the problem from the point of view of KGs in general. The design of a trust layer itself is left as future subsequent work.

KG Summarization & Profiling [23], based on instances or schema, can assist knowledge engineers in evaluating context dimensions already present in the KG. For each property in the role of Qualifier or Referrer, the knowledge engineer would identify which context dimensions it belongs to. It can also help formulate semantic rules interpretations for qualifiers' existence, absence, and possible claim contradictions.

Identifying context dimensions is challenging since any information that characterizes an element can, in principle, be considered context. Knowledge engineers should build on top of an existing KG a set of mappings and rules that makes contextual dimensions explicit about transforming a standard KG into a Contextualized KG (CKG).

References

[1] D. Vrandečić and M. Krötzsch, Wikidata: A Free Collaborative Knowledgebase, *Commun. ACM* **57**(10) (2014), 78–85–. doi:10.1145/2629489.

[2] W. Foundation, Wikidata for authority control, 2022.

[3] S. Cook, The uses of Wikidata for galleries, libraries, archives and museums and its place in the digital humanities, *Comma* **2017**(2) (2019), 117–124. doi:10.3828/comma.2017.2.12.

[4] M. Farda-Sarbas and C. Müller-Birn, Wikidata from a Research Perspective - A Systematic Mapping Study of Wikidata, *CoRR* **abs/1908.11153** (2019). http://arxiv.org/abs/1908.11153.

[5] S. Malyshev, M. Krötzsch, L. González, J. Gonsior and A. Bielefeldt, Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia's Knowledge Graph, in: *The Semantic Web – ISWC 2018*, D. Vrandečić, K. Bontcheva, M.C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L.-A. Kaffee and E. Simperl, eds, Springer International Publishing, Cham, 2018, pp. 376–394. ISBN 978-3-030-00668-6.

[6] D. Schwabe, Trust and Privacy in Knowledge Graphs, in: *Companion of The 2019 World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, S. Amer-Yahia, M. Mahdian, A. Goel, G. Houben, K. Lerman, J.J. McAuley, R. Baeza-Yates and L. Zia, eds, ACM, 2019, pp. 722–728. doi:10.1145/3308560.3317705.

[7] D. Schwabe, C. Laufer and P. Casanovas, Knowledge Graphs: Trust, Privacy, and Transparency from a Legal Governance Approach, *Law in Context. A Socio-legal Journal* **37**(1) (2020), 24–41. doi:10.26826/law-in-context.v37i1.126. https://journals.latrobe.edu.au/index.php/law-in-context/article/view/126.

[8] C. Bizer and R. Cyganiak, Quality-driven information filtering using the WIQA policy framework, *J. Web Semant.* **7**(1) (2009), 1–10. doi:10.1016/j.websem.2008.02.005.

[9] M. van Erp and V. de Boer, A Polyvocal and Contextualised Semantic Web, in: *The Semantic Web - 18th International Conference, ESWC 2021, Virtual Event, June 6-10, 2021, Proceedings*, R. Verborgh, K. Hose, H. Paulheim, P. Champin, M. Maleshkova, Ó. Corcho, P. Ristoski and M. Alam, eds, Lecture Notes in Computer Science, Vol. 12731, Springer, 2021, pp. 506–512. doi:10.1007/978-3-030-77385-4_30.

[10] W. Foundation, Help:Qualifiers, 2023.

[11] K. Shenoy, F. Ilievski, D. Garijo, D. Schwabe and P. Szekely, A study of the quality of Wikidata, *Journal of Web Semantics* **72** (2022), 100679. doi:https://doi.org/10.1016/j.websem.2021.100679. https://www.sciencedirect.com/science/article/pii/S1570826821000536.

[12] A. Piscopo and E. Simperl, What we talk about when we talk about wikidata quality: a literature survey, in: *Proceedings of the 15th International Symposium on Open Collaboration, OpenSym 2019, Skövde, Sweden, August 20-22, 2019*, B. Lundell, J. Gamalielsson, L. Morgan and G. Robles, eds, ACM, 2019, pp. 17:1–17:11. doi:10.1145/3306446.3340822.

[13] L. Saez Trumper D; Pintscher, Identifying Controversial Content in Wikidata, 2022.

[14] T. Yasseri, A. Spoerri, M. Graham and J. Kertész, The most controversial topics in Wikipedia: A multilingual and geographical analysis, *CoRR* **abs/1305.5566** (2013). http://arxiv.org/abs/1305.5566.

[15] A. Chhabra, R. Kaur and S.R.S. Iyengar, Dynamics of Edit War Sequences in Wikipedia, in: *Proceedings of the 16th International Symposium on Open Collaboration*, OpenSym '20, Association for Computing Machinery, New York, NY, USA, 2020. ISBN 9781450387798. doi:10.1145/3412569.3412585.

[16] H.S. Rad and D. Barbosa, Identifying Controversial Articles in Wikipedia: A Comparative Study, in: *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration*, WikiSym '12, Association for Computing Machinery, New York, NY, USA, 2012. ISBN 9781450316057. doi:10.1145/2462932.2462942.

[17] R. Angles, A. Hogan, O. Lassila, C. Rojas, D. Schwabe, P. Szekely and D. Vrgoč, Multilayer Graphs: A Unified Data Model for Graph Databases, in: *Proceedings of the 5th ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, GRADES-NDA '22, Association for Computing Machinery, New York, NY, USA, 2022. ISBN 9781450393843. doi:10.1145/3534540.3534696.

[18] A. Hogan, E. Blomqvist, M. Cochez, C. d'Amato, G. de Melo, C. Gutierrez, S. Kirrane, J.E.L. Gayo, R. Navigli, S. Neumaier, A.N. Ngomo, A. Polleres, S.M. Rashid, A. Rula, L. Schmelzeisen, J.F. Sequeda, S. Staab and A. Zimmermann, Knowledge Graphs, *ACM Comput. Surv.* **54**(4) (2022), 71:1–71:37. doi:10.1145/3447772.

[19] N. Ferranti, A. Polleres, J.F. de Souza and S. Ahmetaj, Formalizing Property Constraints in Wikidata, in: *Proceedings of the 3rd Wikidata Workshop 2022 co-located with the 21st International Semantic Web Conference (ISWC2022), Virtual Event, Hanghzou, China, October 2022*, L. Kaffee, S. Razniewski, G. Amaral and K.S. Alghamdi, eds, CEUR Workshop Proceedings, Vol. 3262, CEUR-WS.org, 2022. http://ceur-ws.org/Vol-3262/paper1.pdf.

[20] A. Piscopo, L. Kaffee, C. Phethean and E. Simperl, Provenance Information in a Collaborative Knowledge Graph: An Evaluation of Wikidata External References, in: *The Semantic Web - ISWC 2017 - 16th International Semantic Web Conference, Vienna, Austria, October 21-25, 2017, Proceedings, Part I*, C. d'Amato, M. Fernández, V.A.M. Tamma, F. Lécué, P. Cudré-Mauroux, J.F. Sequeda, C. Lange and J. Heflin, eds, Lecture Notes in Computer Science, Vol. 10587, Springer, 2017, pp. 542–558. doi:10.1007/978-3-319-68288-4_32.

[21] F. Ilievski, D. Garijo, H. Chalupsky, N.T. Divvala, Y. Yao, C.M. Rogers, R. Li, J. Liu, A. Singh, D. Schwabe and P.A. Szekely, KGTK: A Toolkit for Large Knowledge Graph Manipulation and Analysis, in: *The Semantic Web - ISWC 2020 - 19th International Semantic Web Conference, Athens, Greece, November 2-6, 2020, Proceedings, Part II*, J.Z. Pan, V.A.M. Tamma, C. d'Amato, K. Janowicz, B. Fu, A. Polleres, O. Seneviratne and L. Kagal, eds, Lecture Notes in Computer Science, Vol. 12507, Springer, 2020, pp. 278–293. doi:10.1007/978-3-030-62466-8_18.

[22] F.Å. Nielsen, D. Mietchen and E. Willighagen, Scholia, Scientometrics and Wikidata, in: *The Semantic Web: ESWC 2017 Satellite Events*, E. Blomqvist, K. Hose, H. Paulheim, A. Ławrynowicz, F. Ciravegna and O. Hartig, eds, Springer International Publishing, Cham, 2017, pp. 237–259. ISBN 978-3-319-70407-4.

[23] S. Cebiric, F. Goasdoué, H. Kondylakis, D. Kotzinos, I. Manolescu, G. Troullinou and M. Zneika, Summarizing semantic graphs: a survey, *VLDB J.* **28**(3) (2019), 295–327. doi:10.1007/s00778-018-0528-3.

[24] H. Hanke and D. Knees, A phase-field damage model based on evolving microstructure, *Asymptotic Analysis* **101** (2017), 149–180.

[25] E. Lefever, A hybrid approach to domain-independent taxonomy learning, *Applied Ontology* **11**(3) (2016), 255–278.

[26] P.S. Meltzer, A. Kallioniemi and J.M. Trent, Chromosome alterations in human solid tumors, in: *The Genetic Basis of Human Cancer*, B. Vogelstein and K.W. Kinzler, eds, McGraw-Hill, New York, 2002, pp. 93–113.

[27] P.R. Murray, K.S. Rosenthal, G.S. Kobayashi and M.A. Pfaller, *Medical Microbiology*, 4th edn, Mosby, St. Louis, 2002.

[28] E. Wilson, Active vibration analysis of thin-walled beams, PhD thesis, University of Virginia, 1991.

[29] G. Amaral, A. Piscopo, L. Kaffee, O. Rodrigues and E. Simperl, Assessing the Quality of Sources in Wikidata Across Languages: A Hybrid Approach, *ACM J. Data Inf. Qual.* **13**(4) (2021), 23:1–23:35. doi:10.1145/3484828.