



Contents lists available at ScienceDirect

# Web Semantics: Science, Services and Agents on the World Wide Web

journal homepage: [www.elsevier.com/locate/websem](http://www.elsevier.com/locate/websem)

## A study of the quality of Wikidata

Kartik Shenoy<sup>a</sup>, Filip Ilievski<sup>a,\*</sup>, Daniel Garijo<sup>b</sup>, Daniel Schwabe<sup>c</sup>, Pedro Szekely<sup>a</sup><sup>a</sup> Information Sciences Institute, University of Southern California, USA<sup>b</sup> Ontology Engineering Group, Universidad Politécnica de Madrid, Spain<sup>c</sup> Dept. of Informatics, Pontificia Universidade Católica Rio de Janeiro, Brazil

### ARTICLE INFO

#### Article history:

Received 23 July 2021

Accepted 7 November 2021

Available online 5 December 2021

#### Keywords:

Wikidata

Data quality

Knowledge graphs

Constraints

Crowdsourcing

### ABSTRACT

Wikidata has been increasingly adopted by many communities for a wide variety of applications, which demand high-quality knowledge to deliver successful results. In this paper, we develop a framework to detect and analyze low-quality statements in Wikidata by shedding light on the current practices exercised by the community. We explore three indicators of data quality in Wikidata, based on: (1) community consensus on the currently recorded knowledge, assuming that statements that have been removed and not added back are implicitly agreed to be of low quality; (2) statements that have been deprecated; and (3) constraint violations in the data. We combine these indicators to detect low-quality statements, revealing challenges with duplicate entities, missing triples, violated type rules, and taxonomic distinctions. Our findings complement ongoing efforts by the Wikidata community to improve data quality, aiming to make it easier for users and editors to find and correct mistakes.

© 2022 Published by Elsevier B.V.

### 1. Introduction

Historically, Wikipedia is the best known knowledge base relying on the “wisdom of the crowd” [1] to ensure its quality; setting an example for other popular websites such as Quora<sup>1</sup> and Stack Exchange.<sup>2</sup> Wikidata [2] has been created in a similar manner – editing it is fairly straightforward. Consequently, Wikidata today is a joint creation of tens of thousands of human and bot contributors [3]. The result is a rich set of factual statements that describe claims about entities and events in the real world. New information is entered everyday, resulting in very high growth rates and immediate description of popular world events.<sup>3</sup>

Wikidata aims to allow “plurality of facts” [4], and hence it is important that these facts are described with high-quality statements. We have little understanding of the quality of the knowledge contained in Wikidata. Relatively simple validators can spot syntactic errors, allowing for automatic detection (‘flagging’ or editing) of syntactically anomalous statements [5]. Yet, capturing and fixing semantic information is more challenging. While existing work has proposed an extensive set of quality notions [6], and started to apply statement validation to Wikidata [3,7], to our knowledge, no past work has comprehensively

applied indicators to measure quality of statements in Wikidata as a whole, and provided a vision for improving its quality in the future.

In this paper, we develop a framework to detect and analyze low-quality statements in Wikidata by shedding light on the current practices exercised by the community. In addition, we propose to enhance the quality of Wikidata by automatically flagging potential problematic statements for editors. Our work makes the following contributions:

1. We define three indicators that measure well-understood notions of quality of Wikidata statements, based on: (1) the statement revision history of Wikidata; (2) deprecation of statements; and (3) violations of property constraints defined by the community.
2. We develop an efficient framework that flags potential errors integrating these three indicators of quality. Namely, the community-based indicators find low-quality statements which have been deleted or deprecated throughout the history of Wikidata (since its inception in 2014), while the constraint-based indicator reveals outliers with high constraint violation ratios.
3. We apply our framework to analyze the quality of the entire Wikidata.<sup>4</sup> We report findings on key aspects of quality that affect users and editors, such as low-quality type statements, taxonomical modeling errors, duplicated nodes, and missing statements.

<sup>4</sup> There are 1,149,471,184 statements in the Wikidata dump of December 2020.

\* Corresponding author.

E-mail addresses: [kshenoy@isi.edu](mailto:kshenoy@isi.edu) (K. Shenoy), [ilievski@isi.edu](mailto:ilievski@isi.edu) (F. Ilievski), [daniel.garijo@upm.es](mailto:daniel.garijo@upm.es) (D. Garijo), [dschwabe@inf.puc-rio.br](mailto:dschwabe@inf.puc-rio.br) (D. Schwabe), [pszekely@isi.edu](mailto:pszekely@isi.edu) (P. Szekely).

<sup>1</sup> <https://quora.com/>.

<sup>2</sup> <https://stackexchange.com/>.

<sup>3</sup> In this work, we use the terms *statement*, *claim*, and *fact* interchangeably.

4. We propose recommended actions to interactively support high-quality contributions in the future, as well as to retroactively fix existing issues. By doing so, we complement ongoing efforts by the Wikidata community to improve data quality based on games and suggestions, aiming to make it easier to prevent, find, and correct mistakes.

Our quality indicators evaluate the degree of community consensus on what is acceptable, thus connecting to existing metrics of Wikidata quality, like accuracy, consistency, and veracity [6]. By analyzing statements which have been removed, we reflect on the accuracy of the data. By formulating and analyzing semantic rules (constraints) that statements must satisfy, we provide insights into the well-formedness and consistency of the data. The analysis of the deprecated statements addresses the veracity of claims, by indicating that there was once consensus about their veracity, but this is no longer the case.

We make our code<sup>5</sup> [8] and materials<sup>6</sup> available to facilitate further work on analyzing quality of Wikidata statements. The rest of the paper is structured as follows. Section 2 introduces the three indicators, their formalization, and combination into a joint framework. All of our findings with their supporting analyses are described in Section 3. Recommended actions can be found in Section 4. We relate to prior work on Wikidata quality in Section 5. The paper concludes in Section 6.

## 2. Framework

We seek to measure semantic quality aspects of Wikidata. We devise a framework for detecting low-quality statements in Wikidata, which combines three indicators of quality, based on: (1) community updates; (2) deprecated statements; and (3) property constraints. In this section, we describe each of the quality indicators and we provide details on their formalization into an integrated framework that can analyze the quality of Wikidata. We formalize low quality with  $q = 0$ .

Throughout this section, we use  $S$  to refer to a set of statements. A statement  $(s, p, o, Q) \in S$  refers to the union of an edge subject  $s$ , predicate  $p$ , object  $o$ , and qualifier set  $Q$ . Qualifier sets contain property-value pairs  $(q_{p_i}, q_{v_i}) \in Q$  that further describe the tuple  $(s, p, o)$ , (e.g., with the date or the source of assertion). Such statements are common building blocks of modern hyperrelational Knowledge Graphs (KGs), like Wikidata or YAGO [9].

### 2.1. Quality indicators

**Community-based indicator** We define a community-based indicator of KG quality by considering that the KG statements that have been permanently deleted by the community (i.e., statements deleted at a time point  $t_i$  and not restored in time points  $t_j, j > i$ ) are of low quality. Following the idea of “wisdom of the crowd” [1], we assume that community-based KGs, like Wikidata, are self-correcting over time, i.e., its contributors detect low-quality statements, and either delete or replace them.

However, the set of removed tuples by itself is neither necessary nor sufficient to indicate incorrect statements. A statement might be simply updated with a semantically equivalent one. Object values may be reassigned from one property or class to another, which might be considered more appropriate to express the relationship between the subject and the object. Literals may be updated with a new value that may or may not be semantically different than the original one. The latter case often corresponds

to the adoption of new naming conventions, e.g., replacing the name “Pamela C Rasmussen” with “Pamela C. Rasmussen”. To address these issues, we consider the low-quality ( $q = 0$ ) statements of a dump  $d$  at a time  $t_i$  to be a union of: (1) the removed statements which were not updated ( $R(d_{t_i})$ ), and (2) the removed statements which were updated with a significantly different value ( $U(d_{t_i})$ ). Formally,  $S_c(q = 0, d_{t_i}) = R(d_{t_i}) \cup U(d_{t_i})$ .

**Deprecation-based indicator** Wikidata has a ‘soft’ alternative to deletions: deprecating statements to indicate consensus about the end of their validity. A statement is marked as *deprecated* in two cases: (1) if it has been superseded by another statement, or (2) if it is now known to be wrong, but was once thought correct.<sup>7</sup> For example, the community agreed that Pluto ceased to be a planet since 13th September, 2006 and hence the claim stating that has been deprecated.

Deprecated statements ( $D$ ) are valuable for studying the evolution of Wikidata and the agreement about its statements. However, they are undesired when using Wikidata in applications that require up-to-date information, like entity linking and question answering. Thus, we consider all deprecated statements of a dump  $d$  at a time  $t_i$  to be indicators of low quality, formally:  $S_d(q = 0, d_{t_i}) = D(d_{t_i})$ .

**Constraints-based indicator** The Wikidata community has defined property constraints, i.e., rules that specify how properties should be used.<sup>8</sup> Each property in Wikidata specifies the constraint types that apply to it. Statements expressed with that property can then either conform to the constraint or violate it. We denote the set of all violations in a Wikidata dump  $d$  at a time  $t_i$  with  $V(d_{t_i})$ . Constraints are split in three groups: mandatory, suggested, and normal (i.e., constraints which are neither mandatory, nor suggested). Each constraint type is further specified per property, by stating additional *elements*: property-dependent classes, exceptions, and property paths.

At present, Wikidata defines 30 types of property constraints. Constraints vary in nature, and range from format validation (e.g., correct dates or naming conventions) to ensuring a consistent usage of a property (e.g., making sure that symmetric properties are used in both directions). We provide examples for three key constraint types in Fig. 1: type constraint, value type constraint, and item-requires-statement constraint. The Wikidata type and value type constraints indicate that the domain of a property (or range, respectively) has to conform to one of the listed classes, but specify them further with exceptions and property paths. The item-requires-statement constraint dictates that a Wikidata item with one property should also specify another one. Constraints may also specify *exceptions*. In Fig. 1, the *type* constraint indicates that subjects that have an occupation have to be instances of one of the eight allowed classes, unless the subject is *prescriber* (“person legally empowered to write medical prescriptions”),<sup>9</sup> whereas the *value type* constraint dictates that objects of *occupation* statements have to be either instances or subclasses of one of the six possible classes shown. The *item requires statement* constraint specifies that items which have an *occupation* value must also have an *instance-of* statement. All constraints presented in this figure have a normal status.

The constraint-based indicator considers violations of property constraints in their corresponding statements to be low-quality statements. We denote the set of statements that violate a constraint with  $V$ . The set of low-quality ( $q = 0$ ) statements according to this indicator is:  $S_v(q = 0, d_{t_i}) = V(d_{t_i})$ .

<sup>5</sup> <https://github.com/usc-isi-i2/wd-quality>.

<sup>6</sup> [https://w3id.org/wd\\_quality](https://w3id.org/wd_quality).

<sup>7</sup> <https://www.wikidata.org/wiki/Help:Deprecation>.

<sup>8</sup> [https://www.wikidata.org/wiki/Help:Property\\_constraints\\_portal](https://www.wikidata.org/wiki/Help:Property_constraints_portal).

<sup>9</sup> <https://www.wikidata.org/wiki/Q99393050>.

<b>type constraint</b> exception to constraint relation class	prescriber instance of person narrative entity fictional character animal human robot group of humans group of fictional characters	<b>value type constraint</b> relation class	instance or subclass of occupation profession fictional profession job craft fictional religious position
<b>item requires statement constraint</b> property		<b>instance of</b>	

**Fig. 1.** Example constraints for the property `occupation` (P106): *type* (top-left), *value type* (top-right), and *item requires statement* (bottom). The *type* constraint specifies that subjects that have an `occupation` have to be instances of one of the eight allowed classes, unless the subject is `prescriber`. The *value type* constraint dictates that objects of `occupation` statements have to be either instances or subclasses of one of the six possible classes shown. The *item requires statement* constraint specifies that items which have an `occupation` value must also have an `instance-of` statement. All depicted constraints have a normal status.

## 2.2. Experimental setup

While the three indicators of quality have different foci, each of them identifies a set of low-quality statements, denoted by  $S_c$ ,  $S_d$ , and  $S_v$  in the previous section. In the rest of the paper, we analyze the low-quality statements identified by each indicator. We inspect deprecated and permanently deleted statements in Wikidata, we assess what constraints are violated, and we compare the violations with the deletions. In our experiments, we employ the Knowledge Graph ToolKit (KGTK) [10], which supports flexible and scalable imports of Wikidata, and supports efficient manipulation of large hyperrelational KGs, which is essential for the analysis carried out by our quality framework.

**Community-based indicator:** We collected a dataset of Wikidata statements that have been permanently removed (i.e., removed and not added again) since the first available dump of Wikidata in October, 2014. The dumps of Wikidata are released weekly. We generated this dataset by downloading all available weekly JSON Wikidata dumps from the Internet Archive,<sup>10</sup> resulting in 311 dumps<sup>11</sup>; converting them to the KGTK format; and extracting statements that had been removed between each pair of successive dumps ( $d_{t_i}, d_{t_j}$ ), where  $t_i < t_j$ . We also checked whether statements that have been removed before  $t_i$  were present in the more recent of the two dumps,  $d_{t_j}$ . Formally:

$$A(d_{t_i}, d_{t_j}) = d_{t_j} \setminus d_{t_i}, \text{ with } t_i < t_j$$

$$R(d_{t_i}, d_{t_j}) = d_{t_i} \setminus d_{t_j}, \text{ with } t_i < t_j$$

$$Tr(d_{t_j}) = (Tr(d_{t_i}) \setminus A(d_{t_i}, d_{t_j})) \cup R(d_{t_i}, d_{t_j}), \text{ with } t_i < t_j$$

Here,  $A$  and  $R$  represent the added and deleted statements between  $d_{t_i}$  and  $d_{t_j}$ , respectively. The operator  $\setminus$  represents a difference between two sets,  $\cup$  is the union, and the total removed statements for the 0-th dump is  $Tr(d_{t_0}) = \emptyset$

After obtaining the full set of removed statements, we analyzed how many of the nodes had been redirected to new nodes (i.e., duplicate removal), and computed the distribution of classes and properties being removed. For literals, we investigated whether a value had been entirely removed or updated by computing the similarity between the removed value and the new one. We analyzed the similarity for each literal type separately. For strings, we measured Levenshtein distance between

the removed and the updated text. For dates, we measured the time distance between the removed and the updated date. For quantities, we computed the difference in magnitude between the removed and the new quantity. We consider deleted statements with no update and deleted statements with a notable update to be of low quality (cf. Section 2.1).

**Deprecation-based indicator:** We consider all deprecated statements to be of low quality. Wikidata indicates deprecation through the `rank` qualifier of a statement. We retrieved all statements with a deprecated rank value in the early Jan, 2021 version of Wikidata (the last dump we collected), and we explored their distribution in terms of entities and properties.

**Constraint-based indicator:** We consider statements that violate constraints to be of low quality,  $q = 0$ . We prioritized constraints that are common in Semantic Web research and cover a sufficient number of properties (e.g., `type` and `value type`).

Wikidata has pages with constraint violation reports,<sup>12</sup> which are calculated with an ad-hoc extension of Wikibase.<sup>13</sup> However, it is unclear whether these reports are updated regularly. Given the size of Wikidata, validating its constraints with the Shape Constraint Language (SHACL) or the Shape Expressions Language (ShEx) is computationally prohibitive [11]. Moreover, it is unclear whether these languages can encode exceptions and allowed values in property constraints, and, to the best of our knowledge, there is no available implementation of SHACL/ShEx constraint validators for Wikidata. For this reason, we encoded each constraint type as a KGTK query template. Each template is instantiated once per property, allowing their efficient validation in parallel. Constraint violations for a property are computed in a two-step manner: we first obtain the set of statements that satisfy the constraint for a property, and then we subtract this set from the overall number of statements for that property. We omit constraints defined on external identifier properties, as our aim is to capture semantic and modeling errors in Wikidata.

An example query template is shown in Fig. 2. The query inspects whether the subjects for a property are instances of a class that is allowed by the constraint, or any of its subclasses. If this is the case, or the subject is listed as an exception to the constraint, then the constraint is satisfied for this statement. Notably, for some properties, the `instance of` relation is replaced with a `subclass of` relation. A full example query for one property can be seen in Appendix.

<sup>10</sup> <https://archive.org/search.php?query=wikidata>.

<sup>11</sup> Approximately two years of dumps were missing from Internet Archive, but we were able to retrieve them with the help of contributors from the Wikidata community.

<sup>12</sup> [https://www.wikidata.org/wiki/Wikidata:Database\\_reports/Constraint\\_violations](https://www.wikidata.org/wiki/Wikidata:Database_reports/Constraint_violations).

<sup>13</sup> <https:// Gerrit.wikimedia.org/r/plugins/gitiles/mediawiki/extensions/WikibaseQualityConstraints>.

**Table 1**

Statistics of the constraints: type (Q21503250), value type (Q21510865), item requires statement (Q21503247), inverse (Q21510855), and symmetric (Q21510862). We show the number of properties with (M)andatory, (N)ormal and (S)uggested constraints, and the corresponding number of statements. For the *item requires statement* constraint type  $all \leq M + N + S$ , because properties have multiple constraints with a potentially different status.

Constraint type	#Properties				#Statements	Validation time (in sec.)			
	All	M	N	S		All	Min	Max	Mean
Type	1,456	165	1,280	11	513,424,170	4.95	5231.15	366.16	174.78
Value type	897	106	786	5	182,087,480	11.41	5323.18	352.08	144.15
Item requires statement	527	78	418	97	302,642,146	1.89	2199.57	133.51	58.6
Inverse	110	6	100	4	9,440,925	8.68	646.22	100.69	54.79
Symmetric	38	5	30	3	7,145,197	9.72	527.33	118.44	68.67

kgtk query

```
-i statements instance_of subclass_of_star \
--match statements: (subject)-[id {label:property}]->(object), \
  instance_of: (subject)-[]->(class), \
  subclass_of_star: (class)-[]->(parent)' \
--where 'parent in expected_parents or subject in exceptions' \
--return 'distinct id, subject, property, object' \
```

**Fig. 2.** Example of a query template for the *type* constraint. The set of allowed parent classes for the subject are defined in `expected_parents`, whereas `exceptions` is the set of subjects for which this constraint is not required. If the subject of a statement is an instance of a class in `expected_parents`, or any of its subclasses, then the constraint is satisfied for that statement. The constraint is also satisfied if the subject belongs to the set of `exceptions` defined by the property constraint. Notably, for some properties, the `instance_of` relation is replaced with a `subclass_of` relation.

**Combination of indicators:** Each quality indicator produces a set of statements. We compute the overlap between the deleted statements and the constraint violations as follows. We added all deleted statements to the Wikidata version where we computed the violations, and calculated the number of violations without and with the total removed statements:  $V$  and  $V_{del}$ , respectively. The difference between these two yields the number of violations that were fixed by the removal of the statements ( $V_{fixed}$ ). Formally:

$$V = V(d_{t_i})$$

$$V_{del} = V(d_{t_i} \cup Tr(d_{t_i}))$$

$$V_{fixed} = V_{del} \setminus V$$

### 3. Findings

Our framework indicators result in: (1) a dataset of 76.5M removed statements, describing 26.2M distinct subjects [12]; (2) a dataset of 10M deprecated statements [13]; and (3) a set of correct statements and constraint violations [14], according to the constraint types specified in Table 1. This table shows that most of the property constraints have a normal status, and that the median time to validate a property constraint over Wikidata ranges between 55 and 175 s for the five constraints. This demonstrates the feasibility of our approach to validate Wikidata constraints at scale.

In this section, we highlight the main findings of our analysis by shedding light into complex issues related to KG quality, such as node redundancy, naming conventions, taxonomic distinctions, completeness, accuracy of constraints, and type consistency. We also explore whether constraint violations are getting corrected over time, thereby improving the overall quality of Wikidata. Specifically, we study the following eight research questions:

1. Are entities being deduplicated?
2. Can the community distinguish classes from instances?
3. Are naming conventions needed?

**Table 2**

Distribution of classes in redirected P31 statements. We show 5 classes with the highest number of redirected instances, and 5 classes that have been redirected themselves. The counts and the percentages represent numbers of affected statements. The percentages are relative to total redirected statements, not total statements.

Classes of redirected instances		
Q4167836	Wikimedia category	526,207 (21.38%)
Q5	Human	222,809 (9.05%)
Q4167410	Wikimedia disambiguation page	108,583 (4.41%)
Q13442814	Scholarly article	101,156 (4.11%)
Q7187	Gene	88,231 (3.59%)
Redirected classes		
Q17329259	Encyclopedic article	301,359 (12.25%)
Q4423781	Dictionary entry	53,671 (2.18%)
Q17143521	Village of Poland	51,581 (2.09%)
Q15917122	Rotating variable star	50,642 (2.06%)
Q20900710	Painting	23,482 (0.99%)

4. Are property types and value types respected?
5. Can we detect missing triples?
6. Are constraints correct and complete?
7. What statements get deprecated?
8. Are constraint violations getting fixed?

For each of these questions: (1) we motivate its relevance and impact on Wikidata; (2) we present our findings about its current state; and (3) we provide an in-depth analysis and representative examples. Based on these findings, we provide recommendations about improving the state-of-the-art quality of Wikidata in Section 4.

#### 3.1. Are entities being deduplicated?

Entity linking and deduplication are complex open research challenges in many KGs. Redirects are a common mechanism to deduplicate nodes, and are applied when a user recognizes that two nodes describe the same subject, e.g., `Category:1911 in Morocco` redirects from Q18511155 to Q9404406.<sup>14</sup> Our analysis reveals over 2 million redirected nodes, which affect over 20 million statements (26% of all removed statements). The relatively high number of redirects reflects Wikidata's dynamic nature and the community pursuit for a high-quality, well-integrated graph. It is not known how many duplicate entities currently remain in Wikidata.

21.3 million statements (27.8% of the removed statements) have either a redirected subject or a redirected object. We inspected the property containing the largest number of redirected items, `instance of` (P31), to understand what type of nodes have been redirected. Table 2 (top) shows the five classes with the highest number of redirected instances, which include well-populated classes in Wikidata like human, scholarly article, and

<sup>14</sup> <https://www.wikidata.org/wiki/Help:Redirects>.

**Table 3**  
Community updates of instance-of (P31) and subclass-of (P279).

Before	After	Count	Example
P31	P31	2.85M	(Hardenstein Castle,P31,geographical feature) → (Hardenstein Castle,P31,ruins)
	P279	44k	(laboratory centrifuge,P31,laboratory equipment) → (laboratory centrifuge,P279,laboratory equipment)
	Both	106k	(mystic,P31,person) → (mystic,P31,non-professional work activity) →(mystic,P279,religious)
P279	None	703k	(Clubland Smashed,P31,album)→ <i>none</i>
	P31	444k	(Chemical Markup Language,P279,markup language) → (Chemical Markup Language,P31,markup language)
	P279	33k	(girder bridge,P279,bridge by structural type) → (girder bridge,P279,bridge)
	Both	421k	(barn,P279,building) →(barn,P31,type of farm house) →(barn,P279,agricultural structure) →(barn,P279,appendage)
	None	36.5k	(Categoria:Plantilles d'informació de videojocs,P279, Category:Wikimedia templates) → <i>none</i>

gene. In addition, a portion of the instance of (P31) redirects are due to classes that themselves have been redirected. Table 2 (bottom) shows the five redirected classes with a highest number of member instances, which include encyclopedic article, village of Poland, and rotating variable star.


### 3.2. Can the community distinguish classes from instances?

When adding new instances to Wikidata, contributors must specify descriptive values for the taxonomy relations of instance of (P31) and subclass of (P279). Wikidata's fairly wide ontology (containing millions of classes) and the prior evidence on the difficulty of distinguishing between taxonomic relations in Wikidata [3], raise the question: can the community distinguish classes from instances? Our analysis of removed statements with object properties reveals nearly half a million cases where one of the taxonomic relations has been changed to the other, which point to the fact that the community struggles to decide whether to use instance-of (P31) or subclass-of (P279) to model inheritance in Wikidata.<sup>15</sup>

Drilling down, we see that in 44 thousand cases, the instance of statement was replaced with a subclass of statement. In the case of former P279 edges, the number of taxonomic switches is notably larger: nearly half (444k out of 935k) P279 edges were replaced by a P31 edge only. Illustrative examples in Table 3 indicate that these switches often happen in cases where it is not trivial to distinguish between the two taxonomic relations. For example, the community struggles to specify the membership of laboratory centrifuge as laboratory equipment – a former instance of relation has been replaced with a subclass of one. Conversely, the Chemical Markup Language used to be specified as a subclass of a markup language, but this has been corrected into an instance of relation. In both cases, the updated relation seems more intuitive, which, in line with the “wisdom of the crowd” assumption, would indicate that switches between the two relations largely reflect fixes of prior modeling errors.

### 3.3. Are naming conventions needed?

To our knowledge, Wikidata does not prescribe how to encode strings, though there are guidelines for dates.<sup>16</sup> We performed



**Fig. 3.** Distribution of the Levenshtein distance between old and new string values. The x-axis shows the Levenshtein distance, while the y-axis shows number of statements in each bucket, in terms of millions.

an analysis to investigate the proportion of updates for both strings and dates, in order to study current practices, and possible oscillations between different semantically equivalent values. Our analysis reveals that the community has already performed millions of updates between semantically (nearly) equivalent forms of literals.

In particular, we observe that in the majority of cases (61.5% of all removed dates), the date was replaced with a semantically equivalent date with a different surface form. An example is the year 1964, modified from “000000001964-00-00T00:00:00Z/9” to “1964-00-00T00:00:00Z/9”. When it comes to removed string statements, we observe that 46% of them (14 million) have been replaced with new values. The distribution of the Levenshtein distances between the old and the new string values is shown in Fig. 3. We observe that strings with low Levenshtein distances are typically stylistic updates, e.g., from “Pamela C Rasmussen” to “Pamela C. Rasmussen”. Among the strings with a medium Levenshtein distance (of 10), we see updates which are meant as specifications, such as the update of “Hiroshima EAST BLD” to “Hiroshima East Building”. The strings with a large distance (of 20) are generally different from the original strings, such as the update of “Meredith Boyle Metzger” to “Susan Michaelis”.


<sup>15</sup> [https://www.wikidata.org/wiki/Wikidata:WikiProject\\_Ontology/Problems](https://www.wikidata.org/wiki/Wikidata:WikiProject_Ontology/Problems).

<sup>16</sup> <https://www.wikidata.org/wiki/Help:Dates>.

**Table 4**

Correct (constraint-satisfying) and incorrect (constraint-violating) statements for the five constraint types analyzed in this paper: type (Q21503250), value type (Q21510865), item requires statement (Q21503247), inverse (Q21510855), and symmetric (Q21510862). The violation ratio (VR) is the percentage of incorrect statements in the total set of statements in a given category. We separate the statistics among (M)andatory, (N)ormal and (S)uggested constraints.

Constraint type	Mandatory			Normal			Suggested		
	Correct	Incorrect	VR%	Correct	Incorrect	VR%	Correct	Incorrect	VR%
Type	44.99M	37.67k	0.08	464.71M	3.58M	0.76	85.03k	21.65k	20.29
Value type	11.44M	5.38k	0.03	169.47M	1.11M	0.65	46.15k	512	1.09
I.R.S.	3.98M	767	0.02	272.71M	2.25M	0.82	25.73M	2.24M	8.01
Inverse	6.56k	133	1.99	7.13M	0.21M	2.79	2M	95.35k	4.55
Symmetric	7.43k	42	0.56	6.23M	78.88k	1.25	0.77M	54.22k	6.55



**Fig. 4.** Distribution of Violation Ratios (VRs) for each of the five constraints. Each dot corresponds to a single property. Properties are shown in a descending order according to their VRs. The number of points varies according to constraint usage.

### 3.4. Are property types and value types respected?

Type and value type constraints are similar to the domain and range constraints in Semantic Web languages like OWL, and are covered in resources like YAGO [9] and VerbNet [15]. Many properties in Wikidata have associated type and value type constraints, as shown in Table 1. Have these constraints been respected by the data? We observe that only a small portion of the mandatory constraints, and a much larger portion of the suggested constraints, violate the set constraints. While the violations are largely concentrated around a small set of properties and could in theory be fixed, it is unclear whether this is desired, as the suggested status implies that they might not need to be strictly enforced.

As shown by the violation ratios in Table 4 (rows 1 and 2), only a small portion of the mandatory type and value type constraints are violated (0.08% and 0.03%, respectively). The proportion of violations is larger for normal constraints, which represent the majority (0.76% and 0.65%, respectively). The violation ratio is the highest for the suggested constraints, where as many as 20% of the statements were found to violate type constraints. This might be expected, as the status *suggested* implies less strict semantics than mandatory constraints. This analysis entails that fixing the current type and value type violations would require nearly 44 thousand edits for the mandatory constraints, and 4.7 million edits for the normal and suggested constraints. Fig. 4 shows a Zipfian distribution of the violation ratios for the properties that have type and value type constraint, i.e., most violations are concentrated around a few properties.

### 3.5. Can we detect missing triples?

It is well known that broad-coverage KGs are inherently incomplete [16]. This incompleteness can be partially addressed

through the property constraints: item-requires-statement (IRS), inverse, and symmetric. These constraints point to a missing triple for the same entity, a missing triple with an inverse property, and with a symmetric property, respectively. For example, IRS dictates that entities that have an occupation property must also have a statement with the instance of property. We investigate to which extent these constraints have been followed by the statements in Wikidata. As shown in Table 4, the mandatory constraints for these constraint types reveal nearly a thousand violations, which may indicate missing triples. The situation worsens for normal and suggested constraints, whose enforcement would lead to millions of potentially missing triples. While fixing symmetric and inverse constraints is programmatically trivial, it is unclear whether this is always desired, as the constraint violation may be caused by an incorrect original statement rather than a missing one. For example, if a spouse link exists from entity  $E_1$  to  $E_2$ , but not from  $E_2$  to  $E_1$ , it is impossible to infer automatically whether  $E_1$  and  $E_2$  are spouses (in which case a link from  $E_2$  to  $E_1$  is missing) or not (in which case the link from  $E_1$  to  $E_2$  should be removed).

Table 4 (rows 3–5) illustrates how mandatory IRS and inverse constraints are largely followed (with only 0.02% and 1.9% violations, respectively). As expected, the violation ratios are larger for normal, and largest for suggested constraints, peaking at 8% for the IRS suggested constraints. Table 5 shows examples for properties with highest violation ratios. For instance, the property votes received (P1111) requires other properties like office contested (P541) to be present, which is violated in all 46k cases where it appears. The inverse property for the properties has natural reservoir (P1605) and stepparent (P3448) is missing in nearly all cases, resulting in five thousand violations. The most commonly violated symmetric properties include Sandbox-Lexeme (P5188), together with (P1706), and scheduled service destination (P521), resulting in around 1,500 violations in total.

### 3.6. Are constraints correct and complete?

If the constraints are to be used as a driving force to improve the quality of Wikidata, it is important that they are correct and complete. As shown in Table 4, the majority of the constraints fit the data, which can be seen as an indicator that the constraints are of good quality. Yet, we note that across all constraint types, a small portion of the constraints yields a large portion of violations.

The head of the distribution in Fig. 4 reveals properties whose constraint definitions are outdated. Table 5 lists those property constraints with large (nearly 100%) violation ratios, which may point to discrepancies between the constraints and the underlying data. For example, towards (P5051) expects subjects to be instances of transport stop (Q548662), which is violated for all its 64 instances. 28 of these instances have a type vein (Q9609) (e.g., external jugular vein (Q2512768)), and use the towards property to indicate the direction blood flow of a vein

**Table 5**

Top-3 constraint violations for each constraint type. The violation ratio (VR%) is the percentage of incorrect statements in the total set of statements for a given property.

Constraint	Property	Label	VR%	#Statements
Type	P8138	Located in the statistical territorial entity	100	461
	P5051	Towards	100	64
	P2303	Exception to constraint	100	39
Value type	P5008	On focus list of Wikimedia project	100	331,026
	P6104	Maintained by WikiProject	100	9,764
	P7374	Educational stage	100	32
I.R.S.	P1111	Votes received	100	46,327
	P2302	Property constraint	100	42,211
	P3063	Gestation period	100	549
Inverse	P1605	Has natural reservoir	94.03	201
	P3448	Stepparent	87.97	4,849
	P926	Postsynaptic connection	85.71	7
Symmetric	P5188	Sandbox-Lexeme	100	2
	P1706	Together with	92.85	56
	P521	Scheduled service destination	75.67	1,718

**Table 6**

Top classes (left) and properties (right) in the deprecated statements.

Class	Count	Property	Count
Infrared source (Q67206691)	2,546,256	Instance of (P31)	3,303,204
Star (Q523)	352,194	Proper motion (P2215)	2,236,125
Near-IR source (Q67206785)	60,055	Parallax (P2214)	2,159,860
Astronomical radio source (Q1931185)	43,618	Radial velocity (P2216)	816,191
Galaxy (Q318)	35,768	Distance from Earth (P2583)	461,113

**Table 7**

Violations within the removed statements for each type of constraint.

Constraint	Mandatory	Normal	Suggestion
Type	763k/2.31M (33.04%)	5.3M/34.87M (15.21%)	920/2.29k (40.12%)
Value type	25.4k/211k (12.03%)	198k/8.99M (22.06%)	235/397 (59.19%)
IRS	4.67k/1.28M (0.36%)	192k/4.85M (3.97%)	190k/6.01M (3.17%)
Inverse	37/345 (10.72%)	177k/534k (33.13%)	11.7k/160k (7.27%)
Symmetric	19/307 (6.19%)	7.52M/10.85M (69.37%)	5.05k/37.5k (13.47%)

in the human body (e.g., subclavian vein is oriented towards the brachiocephalic vein). In this case, rather than fixing each statement with a constraint violation manually, one could generalize the constraint, i.e., enhance the type constraint for the `towards` property to allow for instances of `vein`.

### 3.7. What statements get deprecated?

We investigate whether deprecated statements, as a soft alternative to deletions, reveal different behavior compared to removed statements. Among the 10 million statements with deprecated rank in Wikidata, we observe that many belong to the domain of Astronomy. This indicates that the decision between removing and deprecating a statement largely depends on the community and the domain.

Specifically, we found 10,040,256 deprecated statements. The top-5 properties in deprecated statements are shown in Table 6. We observe that all frequently deprecated properties (e.g., proper motion) belong to the domain of Astronomy, and that large portion of the overall deprecations (around 90%) is expressed with these first five properties. In addition, we observe that the deprecated `instance of` statements describe membership of celestial objects, like infrared source, star, and galaxy.

### 3.8. Are constraint violations getting fixed?

Our analysis reveals that Wikidata has millions of deleted statements and constraint violations. Do these two sets overlap? We observe that many of the removed statements violated a constraint, i.e., many of the removals coincide with former violations, thereby improving the quality of Wikidata over time.

Specifically, out of the 2.31 million removed statements for which a mandatory type constraint is defined, a third violated that constraint (Table 7). Most of the former violations correspond to normal and suggested constraints. Overall, we observe that the removed statements fixed millions of constraint violations, including 6 million type violations and 7.5 million symmetric violations.

Notably, constraints could have been fixed or violated through the addition (instead of removal) of statements, which we are not considering in our work and, as such, it is a limitation of our current analysis.

## 4. Recommendations

The knowledge in Wikidata is relatively reliable in comparison to other general-domain KGs [17]. Yet, our analysis reveals a variety of quality aspects of Wikidata that can be improved going forward. Based on our findings, we propose several recommended actions to include in the interactive contributing environment of Wikidata. These recommendations are intended to prevent low-quality statements from being added, as fixing them later might take a large number of edits. The recommendations can complement ongoing efforts by the Wikidata community to improve data quality based on games and suggestions, aiming to make it easier for users and editors to find and correct mistakes.

**Integrate entity linking:** To prevent introducing duplicate nodes, it would be beneficial to provide suggestions for similar entities when these exist. For instance, if the user is introducing a basketball player named “Michael Jordan” who played for Chicago

Bulls, the environment should inform the user that a similar item is already present in Wikidata (with id Q41421).

**Prevent type and value type violations:** When an editor introduces a new entity, its type should be coherent with the type and value type constraints of its properties. When this is not the case, the editor should be warned about a possible violation. Instead of adapting each new statement, the editor may opt to suggest adapting the constraints themselves.

**Introduce format guidelines for strings:** Our analysis showed that a large portion of the literal updates transform the literal between two semantically equivalent forms. We propose having more precise formatting guidelines for strings, aiming to adopt consistent naming conventions. For instance, a guideline for initials of human names may dictate including a letter and a dot (“Pamela C. Rasmussen” rather than “Pamela C Rasmussen”).

**Complement missing data:** Wikidata’s interactive editing environment should propose that the editor makes complete edits, i.e., edits that satisfy the constraints of the affected properties. One way to achieve this would be to suggest that the edits satisfy the constraints of types item-requires-statement, symmetric, and inverse, by either adding the full set of statements that satisfy the constraint, or removing the one violating it. A complementary idea is to include a link prediction method, like HINGE [18] or StarE [19], in order to suggest missing statements based on probabilistic graph patterns.

**Fix statements retroactively:** Given the large number of existing constraint violations, it is important to help the Wikidata community to fix them. One possibility is to leverage Wikidata’s Distributed games<sup>17</sup> approach and create games to help editors efficiently validate and fix the constraints. A good starting point for this are the property constraints with large violation ratios, which were detected through our analysis in Table 5 and Fig. 4. An alternative approach, based on our finding in Section 3.6, is to fix violations automatically with the expectation that after the automatic fixes there will be fewer violations, and it would be more efficient to fix the errors introduced by the automatic fixes than the original ones. Another option is to employ methods that automatically detect errors in KGs [20].

## 5. Related work

The quality of Knowledge Graphs has been studied in existing literature. Chen et al. [21] proposed a framework for evaluating the quality of KGs, consisting of dimensions that quantify their fitness for downstream applications. Similarly, quality metrics from 28 prior papers are surveyed by Piscopo and Simperl [6], and grouped into three dimensions: intrinsic (i.e., accuracy, trustworthiness, and consistency of entities), contextual (i.e., completeness and timeliness of resources), and representation (i.e., understanding, interoperability of entities). Our quality indicators are orthogonal to these metrics, as we consider the consensus of the community for them. In addition, our methods go further by proposing an approach to efficiently evaluate some of the metrics proposed by Piscopo and Simperl [6].

Many of the metrics proposed by Piscopo and Simperl [6] are covered by Färber et al. [17], who compare the quality of modern KGs: Wikidata, YAGO, DBpedia, FreeBase, and OpenCyc. Piscopo and Simperl [3] evaluated the quality of Wikidata from an ontological perspective, using indicators related to quantitative measures of classes and instances (e.g., number of instances and number of properties) and of the richness of classes, relations, and properties (e.g., inheritance richness and class hierarchy depth). Prior work has also investigated whether the quality of a knowledge statement in Wikidata depends on the

engagement of its editor (leader or contributor) [3,22], or the knowledge provenance indicated through the references of a statement [23]. Instead, our work performs a systematic analysis of constraint violations, and assesses whether the removal of statements by the community reduces violations.

Wikidata includes several tools that monitor, analyze, and enforce aspects of quality. The primary sources tool (PST) facilitates a curation workflow for uploading data into Wikidata.<sup>18</sup> The Objective Revision Evaluation Service (ORES) scores revisions automatically, aiming to detect edits which represent vandalism.<sup>19</sup> Recoin (“Relative Completeness Indicator”) [24] extends Wikidata entity pages with information about the relative completeness of the information. Relative completeness is computed by comparing the available information for an entity against other similar entities. Property constraint pages define existing property constraints and report number of violations for a single dump.<sup>20</sup> Our analysis complements the constraint violations reported by Wikidata’s pages, by providing in-depth insights about these violations, and abstracting them into findings and recommendations.<sup>21</sup>

Recently, Wikidata has started moving beyond individual property constraints, representing a higher-level notion of quality in the form of shapes that are meant to provide norms of well-formedness for sub-graphs describing concepts of interest [7], e.g., human.<sup>22</sup> These shapes are collected as Schemas.<sup>23</sup> Each schema defines the desired sub-graph topology describing a given concept, using ShEx shape expressions [7]. Schemas are defined through consensus among specific communities (e.g., molecular biology, software engineering, etc.) interested in standardizing concepts relevant to them.<sup>24</sup> We have not addressed the analysis of Wikidata at this level of abstraction; but the approach described in this work can be naturally extended in this direction. A similar observation can be made about prior work that encodes Wikidata constraints based on the multi-attributed relational structures (MARS) [25], a formal data model for generalized property graphs devised by Marx et al. [26].

Recognizing the complexity of the class and type hierarchy in Wikidata, the authors of YAGO4 hand-crafted a new, principled type hierarchy for Wikidata, specifying constraints in SHACL<sup>25</sup> and OWL [27]; and running scripts to synthesize YAGO by ingesting the data from Wikidata and processing the SHACL expressions. YAGO4 defines constraints on domain and range, disjointness, functionality, and cardinality. The authors report that enforcing these constraints leads to a removal of 132M statements from Wikidata, i.e., 28% of all facts. The constraints defined by YAGO4 overlap partially with the constraints in Wikidata studied in this paper. Subsequent work should compare the findings from validating constraints in YAGO4 and Wikidata, and it should generalize the in-depth analysis done in this paper to other KGs like YAGO4.

Rashid et al. [28] investigated the evolution of 10 classes from DBpedia over 11 of its releases, measuring aspects of: persistence, consistency, and completeness. This effort resembles our

<sup>18</sup> [https://www.wikidata.org/wiki/Wikidata:Primary\\_sources\\_tool#References](https://www.wikidata.org/wiki/Wikidata:Primary_sources_tool#References).

<sup>19</sup> <https://www.wikidata.org/wiki/Wikidata:ORES>.

<sup>20</sup> [https://www.wikidata.org/wiki/Help:Property\\_constraints\\_portal](https://www.wikidata.org/wiki/Help:Property_constraints_portal).

<sup>21</sup> For more information about data quality tools integrated in Wikidata, we refer the reader to: [https://www.wikidata.org/wiki/Wikidata:WikiProject\\_Data\\_Quality](https://www.wikidata.org/wiki/Wikidata:WikiProject_Data_Quality) and <https://docs.google.com/presentation/d/1r7jqzPaHTsXNNqDc2Op1-qSbcFyaFwOSnkEkStp5L3E/edit>.

<sup>22</sup> <https://www.wikidata.org/wiki/EntitySchema:E10>.

<sup>23</sup> [https://www.wikidata.org/wiki/Wikidata:WikiProject\\_Schemas](https://www.wikidata.org/wiki/Wikidata:WikiProject_Schemas).

<sup>24</sup> [https://www.wikidata.org/wiki/Wikidata:Database\\_reports/EntitySchema\\_directory](https://www.wikidata.org/wiki/Wikidata:Database_reports/EntitySchema_directory).

<sup>25</sup> <https://www.w3.org/TR/shacl/>.

<sup>17</sup> <https://wikidata-game.toolforge.org/distributed/#>.



community-based indicator, but it reports analysis over a small data subset, a smaller knowledge graph, and fewer dumps. The goal of Rashid et al. [28] is to identify potential problems in the data processing pipeline, which is orthogonal to our goal of detecting low-quality statements in the knowledge graph itself.

Other work has focused on data validation in KGs. The LOD Laundromat [5] is a large-scale infrastructure that can validate and clean syntactic errors that do not fit the formal specification of RDF, such as bad encoding, undefined URI prefixes, and premature end-of-file markers. Beek et al. [29] devise a toolchain for analyzing the quality of literals in LOD Laundromat's data collection, proposing to automatically improve their value canonization and language tagging. Our work focuses on errors that cannot be detected by methods that check the syntactic validity of typed literals, like illegal dates, and is thus orthogonal to such prior work.

Recent work has assessed quality for specific domains. For instance, Turki et al. [30] report an analysis using ShEx expressions to assess the quality of COVID-19 knowledge in Wikidata. This analysis is more comprehensive than the one reported in our paper, but with a much more limited scope and less generalizable, reflecting the consensus of a specialized community.

Finally, our work relates to efforts that assess the quality of voluntary contributions to large knowledge bases, like Wikipedia [31,32] and Open Street Maps [33,34]. The quality indicators and findings in these works may inspire future research into the quality of large "wisdom of the crowd"-based KGs like Wikidata.

## 6. Conclusions

This paper studies the quality of Wikidata by proposing three quality indicators based on statements that have been (1) permanently removed; (2) deprecated; or (3) violate constraints defined by the community. Our analysis reveals that, while Wikidata is becoming a KG of increasing quality (removing duplicate entities, fixing modeling errors, and removing constraint violations) there is still room for improvement for preventing entity duplication and constraint violations, having consistent guidelines for literals, and completing missing data.

Our findings may complement ongoing efforts by the Wikidata community to improve data quality based on games and suggestions, aiming to make it easier for users to find and correct mistakes. In fact, we are initiating a discussion on how to integrate our methods, findings, and recommendations into Wikidata's infrastructure. Future work will expand our constraint analysis to additional constraint types and properties; investigate the quality of Wikidata over time, its relation to contributor profiles [3]; and will expand our findings by considering additional qualifiers and Refs. Piscopo et al. [23].

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This material is based on research sponsored by Air Force Research Laboratory, USA under agreement number FA8750-20-2-10002. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Research Laboratory or the U.S. Government. Daniel Schwabe was partially supported by grant 309808/2017-0 from CNPq - Brazil.

## Appendix. Sample instantiated query template

The snippet below represents the series of KGTK queries needed to encode the item requires statement constraints (IRS) for property P1321 (place of origin (Switzerland)) in Wikidata.<sup>26</sup> The property has two IRS constraints: (1) each item of the property should be (P31) a human (Q5) and (2) its country of citizenship (P27) should be Switzerland (Q39). There is a single exception to this rule, the person Hans von Flachslanden (Q1583384). The code of the query below is generated automatically with our framework. Comments have been added (with "#") to explain the different parts of the query.

```
kgtk query # Query to retrieve valid entities
-i claims.P1321.tsv # Statements with property P1321
  claims.P31.tsv # Statements with property P31
  claims.P27.tsv # Statements with property P27
--match
'P1321: (node1)-[nodeProp]->(node2),
 P31: (node1)-[]->(node2_P31),
 P27: (node1)-[]->(node2_P27)'
--where 'node2_P31 in ['Q5'] # subject has to be
      # human (Q5)
      and node2_P27 in ['Q39']' # subject should live in
      # Switzerland (Q39)
--return 'distinct nodeProp.id, node1 as 'node1',
  nodeProp.label as 'label',
  node2 as 'node2''
-o claims.P1321.correct_wo_exceptions.tsv
--graph-cache cache.db;
```

```
kgtk ifnotexists # Now we calculate violations
      # of P1321.
-i claims.P1321.tsv
--filter-on claims.P1321.correct_wo_exceptions.tsv
-o claims.P1321.incorrect_wo_exceptions.tsv
```

```
kgtk query # Exclude exceptions, i.e.,
      # Hans von Flachslanden (Q1583384)
-i claims.P1321.incorrect_wo_exceptions.tsv
--match
'(node1)-[]->()' --where 'node1 in ["Q1583384"]'
-o claims.P1321.incorrect_w_exceptions.tsv
--graph-cache cache.db;
```

```
kgtk ifnotexists # Filter exceptions from
      # violations file
-i claims.P1321.incorrect_wo_exceptions.tsv
--filter-on claims.P1321.incorrect_w_exceptions.tsv
-o claims.P1321.incorrect.tsv;
```

```
kgtk cat # Aggregate correct results.
-i claims.P1321.correct_wo_exceptions.tsv
  claims.P1321.incorrect_w_exceptions.tsv
-o claims.P1321.correct.tsv
```

## References

- [1] J. Surowiecki, *The Wisdom of Crowds : Why the Many are Smarter than the Few and how Collective Wisdom Shapes Business, Economies, Societies, and Nations*, Doubleday, New York, 2004.
- [2] D. Vrandečić, M. Krötzsch, Wikidata: a free collaborative knowledgebase, *Commun. ACM* 57 (10) (2014) 78–85.
- [3] A. Piscopo, E. Simperl, Who models the world? Collaborative ontology creation and user roles in wikidata, 2, 2018, <http://dx.doi.org/10.1145/3274410>.
- [4] C. Möller, J. Lehmann, R. Usbeck, Survey on english entity linking on wikidata.
- [5] W. Beek, L. Rietveld, H.R. Bazoobandi, J. Wielemaker, S. Schlobach, *LOD laundromat: a uniform way of publishing other people's dirty data*, in: *International Semantic Web Conference*, Springer, 2014, pp. 213–228.
- [6] A. Piscopo, 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*, 2019, pp. 1–11.

<sup>26</sup> <https://www.wikidata.org/wiki/Property:P1321>.

- [7] K. Thornton, H. Solbrig, G.S. Stupp, J.E.L. Gayo, D. Mietchen, E. Prud'Hommeaux, A. Waagmeester, Using shape expressions (ShEx) to share RDF data models and to guide curation with rigorous validation, in: *European Semantic Web Conference*, Springer, 2019, pp. 606–620.
- [8] K. Shenoy, F. Ilievski, D. Garijo, usc-isi-i2/wd-quality: First notebook release, 2021, <http://dx.doi.org/10.5281/zenodo.5119983>.
- [9] T.P. Tanon, G. Weikum, F. Suchanek, Yago 4: A reason-able knowledge base, in: *European Semantic Web Conference*, Springer, 2020, pp. 583–596.
- [10] F. Ilievski, D. Garijo, H. Chalupsky, N.T. Divvala, Y. Yao, C. Rogers, R. Li, J. Liu, A. Singh, D. Schwabe, P. Szekeley, KGTK: a toolkit for large knowledge graph manipulation and analysis, in: *International Semantic Web Conference*, Springer, Cham, 2020, pp. 278–293.
- [11] I. Boneva, J. Dusart, D.F. Alvarez, J.E.L. Gayo, Shape designer for ShEx and SHACL constraints, in: *ISWC 2019-18th International Semantic Web Conference*, 2019.
- [12] D. Garijo, P. Szekeley, Wikidata removed statements from Oct 2014–Jan 2021, 2021, <http://dx.doi.org/10.5281/zenodo.4686805>.
- [13] K. Shenoy, F. Ilievski, P. Szekeley, Wikidata deprecated statements by Jan 2021, 2021, <http://dx.doi.org/10.5281/zenodo.5120117>.
- [14] K. Shenoy, F. Ilievski, D. Garijo, Constraint violation summaries (Dump: Dec 7th, 2020), 2021, <http://dx.doi.org/10.5281/zenodo.5121276>.
- [15] K.K. Schuler, VerbNet: A broad-coverage, comprehensive verb lexicon, 2005.
- [16] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, W. Zhang, Knowledge vault: A web-scale approach to probabilistic knowledge fusion, in: *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, pp. 601–610.
- [17] M. Färber, F. Bartscherer, C. Menne, A. Rettinger, Linked data quality of dbpedia, freebase, opencyc, wikidata, and yago, *Semantic Web 9 (1) (2018) 77–129*.
- [18] P. Rosso, D. Yang, P. Cudré-Mauroux, Beyond triplets: hyper-relational knowledge graph embedding for link prediction, in: *Proceedings of the Web Conference 2020*, 2020, pp. 1885–1896.
- [19] M. Galkin, P. Trivedi, G. Maheshwari, R. Usbeck, J. Lehmann, Message passing for hyper-relational knowledge graphs, 2020, arXiv preprint [arXiv:2009.10847](https://arxiv.org/abs/2009.10847).
- [20] P. Yao, D. Barbosa, Typing errors in factual knowledge graphs: Severity and possible ways out, 2021, arXiv preprint [arXiv:2102.02307](https://arxiv.org/abs/2102.02307).
- [21] H. Chen, G. Cao, J. Chen, J. Ding, A practical framework for evaluating the quality of knowledge graph, in: *China Conference on Knowledge Graph and Semantic Computing*, Springer, 2019, pp. 111–122.
- [22] A. Piscopo, C. Phethean, E. Simperl, What makes a good collaborative knowledge graph: Group composition and quality in wikidata, in: *SocInfo*, 2017.
- [23] A. Piscopo, L.A. Kaffee, C. Phethean, E. Simperl, Provenance information in a collaborative knowledge graph: an evaluation of Wikidata external references, in: *International Semantic Web Conference*, Springer, 2017, pp. 542–558.
- [24] V. Balaraman, S. Razniewski, W. Nutt, ReCoin: relative completeness in wikidata, in: *Companion Proceedings of the the Web Conference 2018*, 2018, pp. 1787–1792.
- [25] P.F. Patel-Schneider, D. Martin, Wikidata on MARS, 2020, arXiv preprint [arXiv:2008.06599](https://arxiv.org/abs/2008.06599).
- [26] M. Marx, M. Krötzsch, V. Thost, Logic on MARS: Ontologies for generalised property graphs, in: *IJCAI*, 2017, pp. 1188–1194.
- [27] D.L. McGuinness, F. Van Harmelen, et al., OWL web ontology language overview, *W3C Recomm. 10 (10) (2004) 2004*.
- [28] M. Rashid, M. Torchiano, G. Rizzo, N. Mihindukulasooriya, O. Corcho, A quality assessment approach for evolving knowledge bases, *Semant. Web 10 (2) (2019) 349–383*.
- [29] W. Beek, F. Ilievski, J. Debattista, S. Schlobach, J. Wielemaker, Literally better: Analyzing and improving the quality of literals, *Semantic Web 9 (1) (2018) 131–150*.
- [30] H. Turki, D. Jemielniak, M.A.H. Taieb, J.E.L. Gayo, M.B. Aouicha, M. Banat, T. Shafee, E. Prud'Hommeaux, T. Lubiana, D. Das, D. Mietchen, Using logical constraints to validate information in collaborative knowledge graphs: a study of COVID-19 on wikidata, 2020, <http://dx.doi.org/10.5281/zenodo.4445363>.
- [31] D.M. Wilkinson, B.A. Huberman, Cooperation and quality in wikipedia, in: *Proceedings of the 2007 International Symposium on Wikis*, 2007, pp. 157–164.
- [32] N. Raman, N. Sauerberg, J. Fisher, S. Narayan, Classifying Wikipedia article quality with revision history networks, in: *Proceedings of the 16th International Symposium on Open Collaboration*, 2020, pp. 1–7.
- [33] P. Mooney, P. Corcoran, Characteristics of heavily edited objects in OpenStreetMap, *Future Internet 4 (1) (2012) 285–305*.
- [34] C.C. Fonte, V. Antoniou, L. Bastin, J. Estima, J.J. Arsanjani, J.C.L. Bayas, L. See, R. Vatseva, Assessing VGI data quality, in: *Mapping and the Citizen Sensor*, Ubiquity Press, 2017, pp. 137–163.