



Preprint

2026

Public access

This version of the publication is provided by the author(s) and made available in accordance with the copyright holder(s).

Understanding Wikidata Qualifiers: An Analysis and Taxonomy

Falquet, Gilles; Aljalbout, Sahar

How to cite

FALQUET, Gilles, ALJALBOUT, Sahar. Understanding Wikidata Qualifiers: An Analysis and Taxonomy. 2026, p. 19.

This publication URL: <https://archive-ouverte.unige.ch/unige:192809>

© This document is protected by copyright. Please refer to copyright holder(s) for terms of use.

Last deposit update in Archive ouverte UNIGE on 08.04.2026 09:00

Understanding Wikidata Qualifiers: An Analysis and Taxonomy.

Gilles Falquet* ^{1,2} and Sahar Aljalbout† ¹

¹University of Geneva, Switzerland

²ISRI, Geneva, Switzerland

Abstract

This paper presents an in-depth analysis of Wikidata qualifiers, focusing on their semantics and actual usage, with the aim of developing a taxonomy that addresses the challenges of selecting appropriate qualifiers, querying the graph, and making logical inferences. The study evaluates qualifier importance based on frequency and diversity, using a modified Shannon entropy index to account for the "long tail" phenomenon. By analyzing a Wikidata dump, the top 300 qualifiers were selected and categorized into a refined taxonomy that includes contextual, epistemic/uncertainty, structural, and additional qualifiers. The taxonomy aims to guide contributors in creating and querying statements, improve qualifier recommendation systems, and enhance knowledge graph design methodologies. The results show that the taxonomy effectively covers the most important qualifiers and provides a structured approach to understanding and utilizing qualifiers in Wikidata.

1 Introduction

A Wikidata graph is a set of *statements* that describe *entities*. An entity can be either an *item* or a *property* or a *datatype*. A statement asserts that a given entity (the subject) possesses a particular property with a specific value. Statements can be qualified with additional property-value pairs and it should contain at least one reference to an information source. In this paper we use expressions of the form

$$(s, p, v)[q_1 : v_1, \dots, q_n : v_n]$$

to denote a statement in which s is the subject, p and v form the main property-value pair, and the $q_i : v_i$'s are the qualifying property-value pairs.

In [10] the pair (p, v) is called the main snak of the statement and the $q_i : v_i$ pairs are called the qualifier snaks (or qualifiers for short). In this paper, slightly departing from this terminology we will call p the *property* of the statement and q_1, \dots, q_n the *qualifiers* of the statement. Thus, a qualifier is defined as a property used in a $q_i : v_i$ qualifying pair (qualifier snak).

To make the presentation more readable, items and properties will be noted by their English label rather than their identifier (e.g. David Bowman instead of Q4886422). Accordingly, the statement

George C. Scott was married to Colleen Dewhurst from 1960 until their divorce in 1965

*gilles.falquet@unige.ch

†saharaljalbout@gmail.com

will be denoted as:

(George C. Scott, spouse, Colleen Dewhurst)
[start time : 1960, end time : 1965, end cause : divorce] (1)

instead of

(Q182450, P26, Q253916)
[P580 : 1960, P582 : 1965, P1534 : Q93190] (2)

The Wikidata statements form a multi-qualified¹ knowledge graph because each statement may have several values for each qualifier. For example in the statement

(Pare Lorentz, notable work, The Plow That Broke the Plains)
[publication date :10 May 1936, subject has role :film director,
subject has role :screenwriter, subject has role :film editor] (3)

the `subject has role` qualifier has three distinct values.

The introduction of qualifiers makes the Wikidata model more compact than a simple graph model such as RDF 1.1 [4]. In a simple graph model qualifying a statement, e.g. to add a temporal validity interval, requires an explicit reification of the statement or another design pattern, such as Singleton Properties [11], NDFluents [6], Hoganification [7], etc., that in one way or another adds complexity to the model.

Based on its usage by Wikidata contributors, the concept of qualifier clearly fulfills a real need within the community. As of the first of January 2025, 2240 properties were used as qualifiers, appearing in approximately 20% of the statements (331 631 397 statements out of 1 624 567 397 are qualified)²³. However, this abundance presents a spectrum of challenges:

- Selecting the appropriate qualifiers during the creation of new statements becomes a overwhelming task for the contributors.
- When querying the graph, it’s imperative to consider specific qualifiers to formulate queries that yield meaningful results. This becomes particularly crucial with qualifiers that restrict the validity of a statement to some spatial, temporal, or organizational domain, or with qualifiers that express uncertainty about the statement.
- Establishing logical inferences requires elucidating how qualifier values transition from conditions to conclusions. This problem has been deeply studied in [2]

We propose to address the aforementioned issues, by defining a taxonomy of qualifiers based on their actual usage type in the Wikidata graph (shown in Figure 4).

Methodology and structure of the article: To establish a taxonomy of qualifiers, a thorough examination of their semantics and actual usage is essential. Given the extensive number of qualifiers this study focuses on the most important ones. In section 2, we show that a important qualifier is characterized by a high frequency and/or a high diversity of use. We describe how we adapted a diversity index based on Shannon entropy to obtain a diversity index that is suitable for Wikidata qualifiers. To create the qualifier taxonomy we first analyzed a complete Wikidata dump to select the 300 most important qualifiers (in terms of frequency and diversity). Then, starting from an existing taxonomy, introduced in [12], we analyzed the selected qualifiers, checked if they fit into one of the taxonomy’s classes, and updated the taxonomy when necessary. The qualifier analysis was based on their descriptions and their actual use in Wikidata⁴, (the properties of the statements they qualify and their values). This process yielded a taxonomy comprising qualifier categories and subcategories, the characteristics of which are presented in Section 4. In order to characterise this taxonomy, we demonstrate how qualifiers

¹also known as multi-attributed knowledge graph

²script: `src/gen-stats/stmt_count.py`. This notation is used throughout the article to refer to the scripts used to produce the results presented in the article. They can be found on Github at <https://github.com/cui-ke/wikidata-qualifiers> together with their output.

³These figures do not include the “example” statements, as explained in Section 2

⁴For this study we used the 2025-01-01 Wikidata json dump. All the analysis results can be found on Github at <https://github.com/cui-ke/wikidata-qualifiers/tree/main/stats/2025-01-01>

are distributed across the categories and examine cases where apparent classification confusion occurs (Section 5). In Section 6, we explain how this taxonomy can be used to select the right qualifiers when creating new statements and to formulate correct queries over Wikidata. In conclusion, we also show possible applications of this taxonomy beyond those mentioned in the problem statement, in particular in the design of new knowledge graphs.

2 Qualifiers, Frequency and Diversity

In the introduction we defined a qualifier as a property that is used to qualify statements. In the present study we must restrict definition for two reasons:

1. In principle, every property p should be accompanied by an example, which is a statement of the form

$$(p, \text{Wikidata property example}^5, \text{subject})[p : \text{value}]$$

that represents the statement

$$(\text{subject}, p, \text{value}).$$

For example

$$(\text{highest point}, \text{Wikidata property example}, \text{Norway})[\text{highest point} : \text{Galdhøpiggen}]$$

indicates that the statement

$$(\text{Norway}, \text{highest point}, \text{Galdhøpiggen})$$

exemplifies the use of the highest point property.

Therefore every property appears at least once as a qualifier in an example statement. But the example statements cannot be considered as ordinary statements. They are conceptually meta-statements (statements about statements). This is why we won't consider them when analyzing the qualifiers.

2. Some properties may have scope constraints that prevent them from being used as qualifiers. For instance, the statement

$$(\text{instance of}, \text{property constraint}, \text{property scope constraint})[\text{property scope} : \text{as main value}]$$

indicates that the instance of (P31) property can only appear as predicate in a statement and not as qualifier.

Therefore, even though these properties may actually be used to qualify statements, this is considered misuse and they should not be counted as qualifiers⁶.

Taking these two points into account we consider that a qualifier is a property that 1) qualifies at least one statement that is not an example statement and 2) is not disallowed by a property scope constraint. Applying this definition on the 2025-01-01 Wikidata dump we found 1357 properties that are qualifiers and 881 properties used as qualifiers in non-example statements that violate a scope constraint.

To evaluate the importance of a qualifier, we consider both the number of statements it qualifies (its frequency) and the diversity of the different statements it qualifies.

⁵In addition to Wikidata property example (P1855), example statements may be built with Wikidata property example for senses (P5977), Wikidata property example for media (P6685), Wikidata property example for properties (P2271), Wikidata property example for lexemes (P5192), Wikidata property example for forms (P5193)

⁶In the 2025-01-01 dump we found that 880 properties used as qualifiers violated a "property scope" constraint (compare stats/2025-01-01/q-freq.json and stats/2025-01-01/allowed-as-qualifiers.csv)

2.1 Frequency

We define the frequency $F(q)$ of a qualifier q as the number of statements in which q appears at least once.

We performed a frequency analysis on a full Wikidata dump to get an idea of the frequency distribution of the qualifiers.⁷ Figure 1 plots the qualifier frequencies in order of decreasing value. It shows that, as a function of its position in the ranking, the frequency of a qualifier decreases slightly more quickly than a negative exponential. In fact, 41 qualifiers (3.02%) have a frequency higher than 1 000 000 whereas 579 qualifiers (42.7%) have a frequency lower than 100.

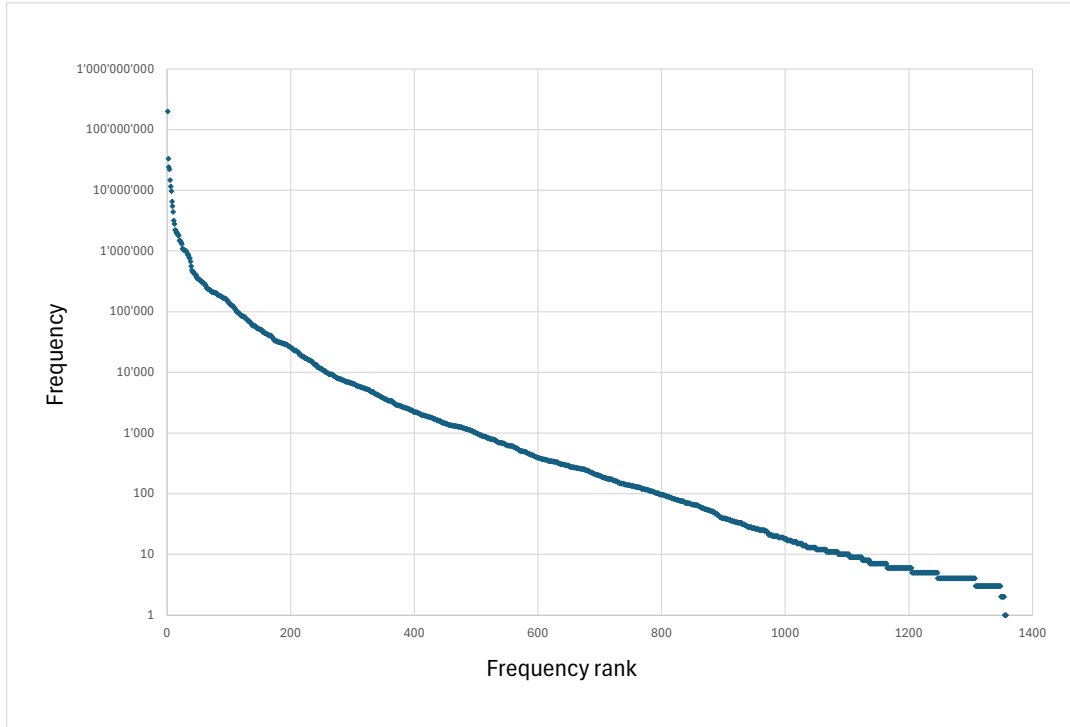


Figure 1: Frequency as a function of the frequency rank for the Wikidata qualifiers

2.2 Diversity

The importance of a qualifier is not necessarily reflected by this frequency alone. A qualifier such as *astronomical filter* (P1227), for example, is extremely frequent (33 145 290 statements qualified), but it is used almost exclusively to qualify statements about the *apparent magnitude* (P1215) of stars (33 145 248 statements). Conversely, the qualifier *valid in place* (P3005) is much less frequent, yet it is used to qualify statements with properties such as *host* (P2975), *number of cases* (P1603), *number of recoveries* (P8010), *instance of* (P31), *box office* (P2142), *number of deaths* (P1120), and 339 others. Therefore the importance of a qualifier involves two dimensions: the number of qualified statements (frequency) and the number of distinct properties appearing in these statements (diversity).

However, the number of distinct properties in the qualified statements (qualified properties) is not directly usable as a diversity indicator. To illustrate this, consider the *catalog* (P972) qualifier. This qualifier is applied to 61 different properties across 24 052 733 statements. However, of these 61 properties, 37 appear fewer than ten times each, meaning that fewer than ten statements are qualified by *catalog* for each of these 37 properties.

Since this “long tail” phenomenon occurs with most of the Wikidata qualifiers it is advisable

⁷script: `src/import/extract_p_q_freq.py`, output: `stats/2025-01-01/p-freq.json`




Figure 2: Frequency/diversity plot of the Wikidata qualifiers where the diversity is defined as the number of distinct properties qualified by a qualifier

to use a well-known diversity index such as the Hill numbers (or effective number of species)⁸, which is commonly used to measure biodiversity and that takes into account the relative frequencies (abundances) of the species.

For this study, we have chosen to use the Hill numbers of order 1. If a dataset contains R types (e.g. species) with relative frequencies p_1, \dots, p_R , the Hill number of order 1 is

$${}^1D := \exp\left(-\sum_{i=1}^R p_i \ln(p_i)\right).$$

This index is the exponential of the Shannon index or entropy.

If the R types in a dataset are equally abundant, with relative frequencies of $1/R$, then the dataset's diversity is R . In other words, a dataset whose diversity index is d has the same diversity as a dataset comprising d equally abundant types.

To apply this formula to a qualifier q let us first define

- the frequency $F(p, q)$ of p qualified by q , as the number of statements with property p that are qualified by q ,
- the properties qualified by q , $P(q)$, as the set of properties p such that $F(p, q) > 0$ (at least one statement with property p is qualified by q).
- the relative frequency $pr(p, q)$ of p for q as

$$pr(p, q) = F(p, q) / \sum_{p' \in P(q)} F(p', q).$$

⁸https://en.wikipedia.org/wiki/Diversity_index

Then the diversity index of a qualifier q is defined as

$$D(q) = \exp \left(- \sum_{p \in P(q)} pr(p, q) \ln(pr(p, q)) \right).$$

The diversity index behaves in the same way as the entropy: If $P(q)$ contains n properties, the index is maximal when the properties have equal relative frequencies of $1/n$. In this case, the index is equal to n . The index is minimal when one property has a relative frequency close to 1, while the others are close to 0; in this case, the index is close to 1. As mentioned above, a diversity index of d means that the qualifier has a diversity that is approximately equal to that of a qualifier which qualifies d properties with the same relative frequencies.

For example, consider the qualifier **genomic assembly** (Q659), whose property frequencies are shown in Table 1, its Shannon index $D(Q659)$ is 3.26. This effectively reflects the fact that it is essentially used with the properties **genomic end**, **genomic start**, **strand orientation**, and **chromosome**. Its use with the other four properties is accidental or incorrect.

| p | Property name | $F(p, Q656)$ | $pr(p, Q656)$ |
|-------|---------------------------|--------------|---------------|
| P645 | genomic end | 514 956 | 0.3457874454 |
| P644 | genomic start | 514 955 | 0.3457867739 |
| P2548 | strand orientation | 422 711 | 0.2838459147 |
| P1057 | chromosome | 36 553 | 0.0245449484 |
| P2043 | length | 46 | 0.0000308885 |
| P1855 | Wikidata property example | 3 | 0.0000020144 |
| P3331 | HGVS nomenclatur | 2 | 0.0000013429 |
| P4844 | research intervention | 1 | 0.0000006714 |

Table 1: Frequencies and relative frequencies of the properties in $P(Q656)$ (**genomic assembly**)

Nevertheless, this index is still not entirely satisfactory, as some properties tend to overshadow the others, drastically reducing the diversity index. For instance, the **series ordinal** qualifier has the frequencies shown in Table 2. Due to the very high frequency of **author name string**, the diversity index is only 2.03. This does not reflect the fact that this qualifier is very frequently used with many other properties. Notably, it qualifies 99% of statements with the property **expressed in** and 94% of statements with the property **author**.

| p | Property name | $F(p, P1545)$ |
|-------|----------------------------------|---------------|
| P2093 | author name string | 148 015 200 |
| P50 | author | 31 769 610 |
| P2860 | cites work | 14 782 837 |
| P5572 | expressed in | 1 056 827 |
| P179 | part of the series | 745 165 |
| P735 | given name | 739 798 |
| P527 | has part(s) | 352 459 |
| P793 | significant event | 210 584 |
| P4908 | season | 170 291 |
| P5753 | ideographic description sequence | 90 169 |
| P361 | part of | 83 594 |
| P26 | spouse | 79 612 |
| P734 | family name | 68 516 |
| P658 | tracklist | 43 553 |
| P710 | participant | 37 559 |
| etc | | |

Table 2: Frequencies of the properties qualified by **series ordinal** (P1545)

Thus we opted to replace the frequency $F(p, q)$ of a property p for a qualifier q by the proportional frequency

$$PF(p, q) = F(p, q)/GF(p)$$

where $GF(p)$, the global frequency of p , is the number of statements having property p . $PF(p, q)$ is the proportion of the statements having property p that are qualified by q .

| p | Property name | $GF(p)$ | $F(p, P1545)$ | $PF(p, P1545)$ |
|-------|----------------------------------|-------------|---------------|----------------|
| P2093 | author name string | 148 810 279 | 148 015 200 | 0.995 |
| P50 | author | 33 552 582 | 31 769 610 | 0.947 |
| P2860 | cites work | 303 125 156 | 14 782 837 | 0.049 |
| P5572 | expressed in | 1 058 470 | 1 056 827 | 0.998 |
| P179 | part of the series | 1 003 261 | 745 165 | 0.743 |
| P735 | given name | 7 901 359 | 739 798 | 0.094 |
| P527 | has part(s) | 2 405 547 | 352 459 | 0.147 |
| P793 | significant event | 1 129 126 | 210 584 | 0.187 |
| P4908 | season | 172 152 | 170 291 | 0.989 |
| P5753 | ideographic description sequence | 90 169 | 90 169 | 1.000 |
| P361 | part of | 5 104 617 | 83 594 | 0.016 |
| P26 | spouse | 830 942 | 79 612 | 0.096 |
| P734 | family name | 5 346 143 | 68 516 | 0.013 |
| P658 | tracklist | 53 797 | 43 553 | 0.810 |
| P710 | participant | 860 972 | 37 559 | 0.044 |
| etc. | ... | ... | ... | ... |

Table 3: Global frequency (GF), frequency (F), and proportional frequencies (PF) of the properties qualified by series ordinal (P1545)

For the above example, the proportional frequencies are shown in Table 3. Using the proportional frequencies instead of the frequencies to compute the diversity index yields a value of 57.64, which better reflects the usage of this qualifier. Interestingly, for genomic assembly, the new index remains at 3.41, close to the original value. This is desirable because the last four properties are incorrectly or accidentally qualified.

Figure 3 presents the position of each qualifier based on its frequency and 1D proportional diversity. It graphically shows that there is no correlation between these two variables.




Figure 3: Frequency and 1D proportional diversity of the Wikidata qualifiers

3 Categorization Approach for Qualifiers

The taxonomy construction is based on an analysis of the most important qualifiers, in terms of frequency and diversity, from the the Wikidata graph. Frequency and diversity computations were carried out by extracting three dictionaries⁹ from the 2025-01-01 Wikidata dump:

$\mathbf{P} \rightarrow \mathbf{F}$ associates every property with its frequency (the number of statements with this property)

$\mathbf{Q} \rightarrow \mathbf{F}$ associates every qualifier with its frequency (the number of times it appears in a statement)

$\mathbf{Q} \rightarrow (\mathbf{P} \rightarrow \mathbf{F})$ associates with each qualifier q a dictionary that associates with each property p the value of $F(p, q)$ (the number of times the qualifier q appears in a statement with property p)

From these data we computed an importance score for each qualifier¹⁰, which is defined as

$$score = frequency \times diversity$$

Since the frequency and diversity values are expressed in different units (frequency is a number of statements, diversity is a number of properties) and since their value ranges differ greatly, it is preferable to combine them by multiplication rather than by weighted addition. This is particularly important to obtain a ranking that is independent of the normalization functions, as explained in, for example, [15].

We selected the top 300 qualifiers as a basis to build a taxonomy of qualifiers. These qualifiers cover 99.6% of the 96 281 446 qualifications present in the graph, i.e. 99.6% of the qualifier-value pairs have a qualifier belonging to the selected set.

Each selected qualifier has been analyzed based on its textual description and on the properties it most frequently qualifies. As the taxonomy aims to reflect the actual use of qualifiers, qualifier analysis could not be based solely on the qualifier descriptions. This is because qualifiers are properties that are used in a qualifying role. However, the description of many properties refers only to their role as statement properties; their role as qualifiers is not described. In fact, only 148 of the 300 selected qualifiers have a description containing a descriptive string, usage instructions, or examples that explicitly mention their role as a qualifier. Furthermore, the textual descriptions are generally limited to one sentence and do not specify the intended use of the qualifiers precisely.

To determine how qualifiers are actually used, we consulted the $\mathbf{Q} \rightarrow (\mathbf{P} \rightarrow \mathbf{F})$ dictionary to identify the most frequently qualified properties. This generally provides a good indication of the actual (operational) meaning of the qualifier for the Wikidata contributors. For instance, the properties most frequently qualified by *language of work or name* (P407) are (in decreasing order of frequency) *official website* (P856), *described at URL* (P973), *work available at URL* (P953), *Fandom article ID* (P6262), *pronunciation audio* (P443), *has edition or translation* (P747), *YouTube video ID* (P1651), *voice actor* (P725), *curriculum vitae URL* (P8214), etc. Almost all of them have values that are URLs or other external identifiers or other resources that are not Wikidata items¹¹. This means that the *language of work or name* qualifier does not describe the statement itself but the external resource designated by the statement's value. We will see in the next section that this kind of qualifiers fall into the External Entity Description category.

If this was insufficient to determine the meaning of the qualifier, we queried the Wikidata graph to find the values most frequently associated with it. This was necessary for qualifiers such as *sourcing circumstances* that exhibit high diversity. In this case the list of qualified properties is very heterogeneous and does not help in finding the meaning of the qualifier.

⁹script: `src/import/extract_p_q_freq`, output: `stats/2025-01-01/pfreq.json`, `stats/2025-01-01/q_freq.json`, `stats/2025-01-01/p-q_freq.json`

¹⁰script: `src/import/diversity_index.py`, output: `stats/2025-01-01/q-diversity-score.csv`

¹¹By the way, *has edition or translation* has a value which is a Wikidata item that should be described by a *language of work or name* property. Therefore the qualification by *language of work or name* is modeling mistake (redundancy)

The categorisation of qualifiers started with Patel-Schneider’s classification [13] that distinguishes between *contextual* and *additional* qualifiers. Contextual qualifiers restrict the validity of a claim made in a statement to a particular context. For example:

(IBM, official website, <https://www.ibm.com/de-de/>)
[applies to jurisdiction : Germany] (4)

means that *<https://www.ibm.com/de-de/> is the official website of IBM* only in German jurisdiction. Additional qualifiers, on the other hand, have no influence on the semantics of the statement, they provide supplementary information, as in .

(common dandelion, taxon name, *Taraxacum officinale*)
[taxon author : Friedrich Heinrich Wiggers] (5)

We also considered other ways to characterize statements, such as causality, uncertainty, confidence, provenance, and context description. These have been amply studied in the knowledge representation literature. For an introduction to some of these aspects, see, for instance, [14] .

4 Results

The classification we propose is a refinement and an extension of the classification by Patel-Schneider [12]. By studying the most important qualifiers, as explained in the previous section, we found that some non-contextual qualifiers are not additional. These qualifiers cannot be removed without altering the meaning of the statement. These qualifiers either carry epistemic information, generally defining the uncertainty (or imprecision) of a statement), or participate in specifying a structured value or defining a data structure (lists, categories) or constraint. We also found that the contextual and additional qualifier category can be refined into several non-overlapping subcategories. This resulted in the taxonomy shown in Figure 4.

In the rest of this section we describe the characteristics of each qualifier category.

4.1 Context Qualifiers (aka Validity)

These qualifiers limit the validity of a statement to a specified context. Without this restriction, the statement would be universally valid for the given subject and object.

A context can be defined along several independent dimensions, such as time, space, or an abstract region (e.g. a jurisdiction or an organization), In each dimension, the validity qualifiers define a subset of all possible values, whether temporal, spatial, etc. Therefore, the validity of a statement is the Cartesian product of the validity subsets for all the considered dimensions.

Temporal context

A temporal context is usually defined using qualifiers such as *point in time*, *start time*, *end time*, or *valid in period*. This is generally a time interval, which can be reduced to an instant, as in

(Peoria, population, 190985)
[point in time : 1 April 2020, ...]

The absence of a temporal context specification can have various meanings. For example, if the statement is about something that cannot change over time, such as the location of the 1969 Woodstock music festival, the notion of temporal validity does not apply. In other words, the statement is always valid. If the statement’s subject is an enduring item (something whose properties can change over time), the absence of a temporal context specification usually means that the statement is valid for the subject’s entire lifetime up to the present day., e.g.

(Eiffel tower, number of floor, 3).

However, it may also mean that the statement is currently valid, but was not always valid, or may not always be valid, e.g.

(Aletsch Glacier, thickness, 900 metre)

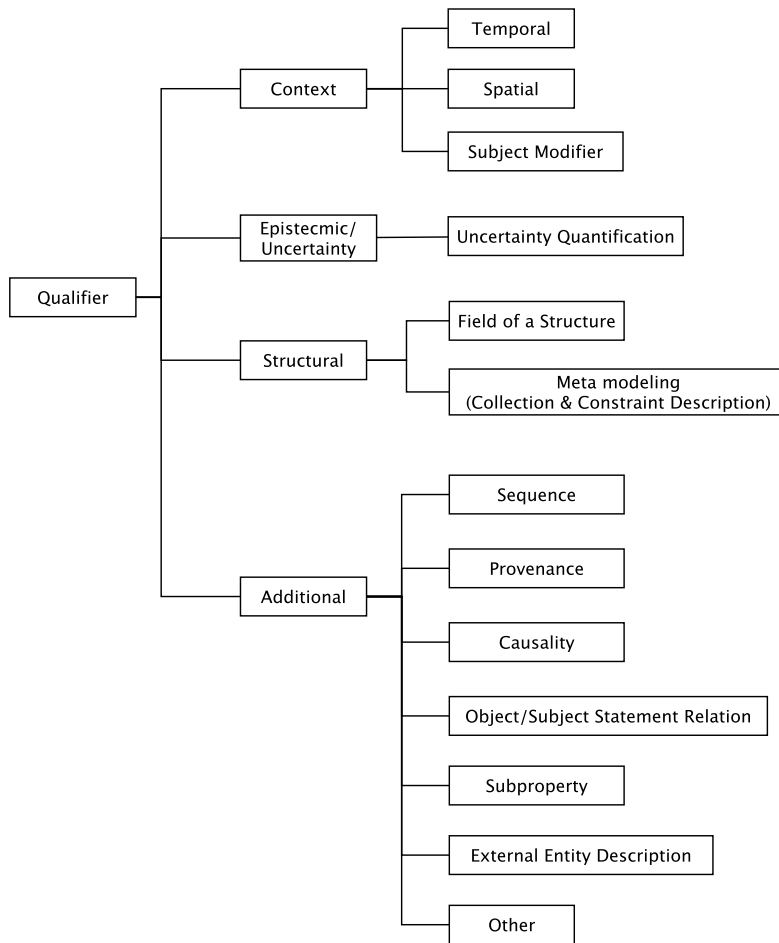


Figure 4: Qualifier taxonomy

Spatial context

Qualifiers such as country or valid in place define the spatial validity of a statement. See for example statement 4.

Subject modifiers

These qualifiers specify the validity of a statement by either extending or restricting the subject to which it applies. Typical qualifiers in this category are: applies to part, including, and excluding.

(Dance of Time, creator, Jean-Baptiste Lepaute)
 [applies to part: clock, ...]

Remark. The qualifier description can sometimes be misleading. For example, the description of *catalog* contains “as a qualifier of P528 – catalog for which the ‘catalog code’ is valid”. This suggest that it is a context qualifier, as it defines the validity domain of a code. However, *catalog* can hardly be considered a context qualifier because, unlike other context qualifiers, it has a low diversity (4.36) , and mostly qualifies *catalog code* (P528) and only a few other identification number properties. It is much more natural to consider that *catalog* is a field of a structured value consisting of a pair (*catalog code*, *catalog*).

4.2 Epistemic/Uncertainty

A few qualifiers, in particular **sourcing circumstances** (P1480) and **nature of statement** (P5102) relate to our knowledge of a statement or the extent to which it has been validated. They can indicate that the statement's value is imprecise, as in:

(Rome, inception, 21 April 753 BCE)[**sourcing circumstances**: circa],

or uncertain, as in:

(Cyrillic script, creator, Constantine of Preslav)
[**sourcing circumstances**: hypothesis].

These qualifiers can also express a level of confidence in the statement, with values such as **official**, **unofficial**, **disputed**, **cannot be confirmed by other sources**, or even the falsity of a statement

Some of these qualifiers are intended to quantify uncertainty in a value.

(Plato, date of birth, 420s BCE)
[**earliest date**: 428 BCE, **latest date**: 427 BCE]

The most frequently used among these qualifiers are **latest date**, **earliest date**, **latest start date**, **earliest end date**, **latest end date**. Note that these qualifiers do not define a temporal validity context, instead, they indicate that the true value of a **date**, **start date**, or **end date** lies somewhere within the given interval.

4.3 Structural

Structural qualifiers participate in the definition of data structures, at either the data type level (metamodeling) or at the value level (a field within a structured value).

Metamodeling

These qualifiers are used to enrich the Wikidata model by defining categories, list types, and constraints.

The Wikidata pattern to define a category *c* that contains entities of type *t* with a specific value *v* for a property *p* consists of a statement (*c*, **category contains**, *t*)[*p*: *v*]. Here, *p* appears as a qualifier, although its main use is as the main property of a statement. For example

(Category:Grade I listed buildings in Bedfordshire,
category contains, architectural structure)
[**located in the administrative territorial entity**: Bedfordshire]

Properties that are often used to define categories include : **occupation**, **performer**, **sex or gender**, **member of sports team**, **located in the administrative territorial entity**, **educated at**, **place of death**, **date of birth**, **position held**, **director**, **country**.

Similarly, a *list* *l* whose elements are of type *t* and have a value *v* for a property *p* is defined by a statement of the form (*l*, **is a list of**, *t*)[*p*: *v*].

A *constraint* on a property *p* is defined by a set of statements of the form (*p*, **property constraint**, *constraint type*)[*q*₁: *v*₁, ..., *q*_{*n*}: *v*_{*n*}]. For example

(located on linear feature, **property constraint**, value-type constraint)
[**class**: railway line, **class**: road, ...]

Typical qualifiers that appear in the definition of constraints are : **class**, **relation**, **item of property constraint**, **property scope exception to constraint**, and **constraint clarification**.

Field of a structured value or subject

These qualifiers are used when the object of a statement is an entity that is not a Wikidata item and must be described by several fields. If one does not want to create a new item in this situation, the technique is to create a statement in which one of the fields is the object and the others are qualifiers, as in the following examples:

- The apparent magnitude of an astronomical object a is a pair (b, f) where b is the brightness and f is the wavelength (corresponding to a filter). It is represented as $(a, \text{apparent magnitude}, b)[\text{astronomical filter}: f]$
- The location of a gene in the genetic material is a number and a chromosome identifier: $(g, \text{genomic start}, n)[\text{chromosome}: c]$
- The catalog code of some entity is a string together with a catalog identifier: $(e, \text{catalog code}, s)[\text{catalog}: c]$
- The boiling point of a substance is a temperature and a pressure: $(s, \text{boiling point}, t)[\text{under pressure}: p]$

These qualifiers generally have a low to very low diversity index (< 6), i.e. they are highly specific to a property. For example, the **astronomical filter** qualifier applies almost exclusively to **apparent magnitude**. Those with a higher index apply to several properties belonging to a common domain. For example, **temperature** (diversity = 24.34) qualifies several physical values (density, vapor pressure, solubility, etc.)

The properties they qualify generally do not make sense without the qualifiers. For instance, the position of a gene is an integer number (**genomic start**) that must be qualified by the chromosome on which the gene is located.

4.4 Additional information

These qualifiers provide additional information, without altering the meaning of the statement. The subcategories that we have identified are: sequence, provenance, causality, the relation between object/subject and statement, sub-property, external entity descriptions.

Sequence

Sequence qualifiers typically indicate the absolute or relative position of the subject of a statement among the entities that are the subject of a statements with the same property and value. For example, in

$(\text{Jimmy Carter}, \text{position held}, \text{President of the USA})[\text{series ordinal} : 39]$

the subject (Jimmy Carter is the 39th entity that has the value President of the USA for the property position held. Sequence qualifiers can also apply to the value of a statement. In

$(\text{Charlie Chaplin}, \text{child}, \text{Geraldine Chaplin})[\text{series ordinal} : 4]$

the value (Geraldine Chaplin) is the 4th entity that is a value of the property child for the item Charlie Chaplin.

The most frequently used sequence qualifiers are **series ordinal** (P1545), **follows P**(155), followed by (P156), **replaces** (P1365), **replaced by** (P1366), **candidate number** (P4243), **candidate position** (P10777).

Provenance

In Wikidata, the provenance (source) of a statement is indicated by its *references*. As every statement should be based on at least one reference, Wikidata has its own specific mechanism for handling references, hence there is no reference qualifier. However, some qualifiers do provide information about the provenance of a statement.

object named as (P1932) and **subject named as** (P1810) indicate how the statement's object and subject are name in the original source ; **determination method or standard** (P459) specify the process or method or standard used to produce the value ; **reason for deprecated rank** (P2241) and **reason for preferred rank** (P7452) explain the ranking of a statement ; **statement supported by** (P3680) provide the source of a value, as in

$(\text{Albany}, \text{per capita income}, 31,800 \text{ New Zealand dollar})$
[statement supported by: 2018 New Zealand census,
point in time: 2018]

Causality

These qualifiers describe what caused the statement to become true or cease to be true, or the consequences of the statement (i.e. the statement describes the cause of something else). The most important causality qualifiers are : **end cause**, **has cause**, **has immediate cause**, **is immediate cause of**, **content descriptor** (typically provides the reason/cause for a rating), **has effect**, **cause of destruction**, **cause of death**.

Object/Subject-Statement relation

These qualifiers provides a relation between the object (resp. subject) and some entity that holds in the context of the statement. This is typically the role played by the subject or object in the context of the statement, as in

(2001 A Space Odyssey, cast member, Keir Dullea)[character role: David Bowman]

for the object role, or in

(Falk Kalamorz, employer, Institute for Crop and Food Research)
[subject has role : scientist]

for the subject role.

Typical qualifiers in this category are: **object of statement has role** (P2868), **subject has role** (P2868), **character role** (P453), **position held** (P39), **position played on team / speciality** (P413).

Sub-property or sub-value (Relation or Value Refinement)

Qualifiers such as **criteria used** or **mapping relation type** make the meaning of a statement more precise by providing additional information about the relation between the item and the value. For instance, in the statement

(Řepka, number of houses, 13)
[point in time : 1980,
criteria used : number of permanently inhabited houses]

the **criteria used** qualifier indicates that *number of houses* is in fact the *number of permanently inhabited houses*.

Similarly, some qualifiers indicate the specific part of the value that is relevant for the statement. In

(Promotion of science and research, main regulatory text, Fiscal Code of Germany)
[section, verse, paragraph, or clause : §52 II Satz 1 Nr. 1]

the qualifier specifies the precise entry in the Fiscal Code of Germany that regulates the Promotion of science and research.

Theses qualifiers provide a view of the statement at a finer granularity level.

External object description

When the object of a statement is the identifier (URL, DOI, ISBN, catalog ID, version number, ...) of an entity that is not described in the Wikidata graph (it is not a Wikidata item), these qualifiers are used to describe this entity. In fact, they act as predicates with the external object as subject. In the following statement the **language of work** qualifier describes the entity (a document) whose URL is the statement's object :

(The Treasure of the City of Ladies, full work available at URL,
<https://www.gutenberg.org/files/26608/26608-h/26608-h.htm>)
[language of work or name : French]

In the statement

(Terraform, software version identifier, 1.7.2)[publication date : 31 January 2024]

the publication date is that of version 1.7.2 of the Terraform software, which is not a Wikidata item (there is only one Wikidata item that represents the software as a whole).

Typical external object description qualifiers are : online access status, author last names, data size, copyright license, distribution format.

5 Characteristics of the taxonomy

5.1 Coverage

Figure 5 shows how the top 50 qualifiers¹² (representing 93% of the qualifications in Wikidata) are distributed in the different taxonomy categories. There is only one category Meta which contains no qualifiers. In the distribution of the top 300 qualifiers shown on Figure 6, the Other Additional catches more than half of the qualifiers. This is because this category contains mainly domain-specific and less important qualifiers. In fact, qualifiers in this class have lower frequencies, as shown in Fig 2, and their average diversity is very low compared to the others, as shown in Table 3 .




Figure 5: Number of qualifiers in each category for the top 50 qualifiers

5.2 Multiple classification / confusion

As already mentioned, the classification of a qualifier is not always immediate and may require the observation of its actual use. However, most of the time the qualifier clearly belongs to only one class.

For the qualifiers that fall into more than one category, we identified some possible causes

- The intended use of a property as a qualifier is not well defined or not defined at all. For instance, the country qualifier can

¹²data: `classification/Qualifier_Analysis_2025-01-01.xlsx`




Figure 6: Number of qualifiers in each category for the top 300 qualifiers

1. indicate the spatial validity of a statement, as for `headquarters location` or `distributed byor`
2. add information about the statement's object when used with `place of birth` or `location`

The latter case is generally a misuse. For a statement $(x \text{ place of birth } y)[\text{country: } c]$ the country of place y should normally be found in a statement $(y \text{ country } c)$. Therefore using a `country` qualifier is redundant. Nevertheless, it occurs for 16 191 items in Wikidata, as of 2025-01-29. This can be caused by the absence of description for `country` as a qualifier and also because `country` is in fact in the list of allowed qualifiers for `place of birth`.

- In the `publisher` case, there is clearly an ambiguity in the qualifier definition. For the `population` property, this qualifier indicates who published this figure, it is a provenance indication. For the `described at url` property, this qualifier refers to the publisher of the web page, not to the author of the statement. It is an external entity description, not a provenance. The same occurs with `point in time`.

6 Usefulness of the taxonomy

So far, we have shown that the proposed taxonomy adequately encompasses all the selected qualifiers, meaning each qualifier falls within a specific category. Of course the taxonomy could be refined further to identify subclasses within the Additional/Other class. However, this would mostly affect less important qualifiers. We also showed that the classes are clearly defined and non-overlapping, although in practice, some qualifiers belong to several classes due to misuse or ambiguous definitions of the qualifier). We must now return to the challenges stated in the introduction and demonstrate how this taxonomy addresses them.

6.1 Selecting the right qualifiers.

Organising qualifiers in a taxonomy helps users explore the qualifier space and find those that apply to their statements. This is true of the proposed taxonomy.




Figure 7: Sum of the frequencies of the qualifiers in each subcategory

1. The categories are defined by clear criteria. The definition of each category provides a way to determine whether a qualifier belongs to the category.
2. The categories correspond to well-known notions that have been extensively studied in the field of data and knowledge representation, such as context (time, space, etc.); uncertainty; causality; provenance; ordering (sequences); and structured objects. Therefore, they are likely to provide a natural set of qualifying dimensions with which the user is familiar and which correspond to their intentions.

6.2 Formulating correct queries and inferences (through qualifier abstraction)

The abundance of qualifiers, makes it difficult to formulate queries or inference rules that take them into account correctly and completely. The taxonomy of qualifiers can help to solve this problem by providing a higher level of abstraction. To achieve this, each category is associated with a data type and each statement is associated with a value for each relevant category. The value for a category is formed from the set of qualifier-value pairs in which the qualifier is in that category.. For instance, the statement

$(s, spouse, v)[start\ time : 1975, has\ cause : marriage, end\ cause : divorce, end\ time : 1982]$

has

- a *temporal context* value [start time : 1975, end time : 1982]
- a *causality* value [has cause : marriage, end cause : divorce]

We can then define high level operations for these data types and use these in queries and inference rules rather than referring directly to the qualifiers.. For instance, consider an inference rule of the form

$$(s_1, p_1, v_1)[T_1] \wedge (s_2, p_2, v_2)[T_2] \rightarrow (s_3, p_3, v_3)[T_3]$$

| Category | Subcategory | Average Diversity |
|-----------------------|-----------------------------------|-------------------|
| Context | Temporal | 106.09 |
| | Spatial | 17.90 |
| | Subject Modifier | 65.00 |
| Epistemic/Uncertainty | Epistemic | 34.14 |
| | Uncertainty Quantification | 14.36 |
| Structural | Structured Value | 7.54 |
| | Meta modeling | 5.64 |
| Additional | Sequence | 14.82 |
| | Provenance | 90.40 |
| | Causality | 20.97 |
| | Object/Subject Statement Relation | 27.78 |
| | Sub-Property or Value | 14.49 |
| | External Object Description | 16.60 |
| | Other Additional | 6.13 |

Table 4: Average diversity of the qualifiers for each (sub)category

where the T_i 's are the temporal context values of the statement. To obtain a correct rule one can use a predicate $\text{intersects}(T_1, T_2)$, which checks if the two statements have time validity that overlap, and an operation $\text{intersection}(T_1, T_2)$ that computes the intersection of temporal validity intervals. Thus we obtain a corrected rule

$$(s_1, p_1, v_1)[T_1] \wedge (s_2, p_2, v_2)[T_2] \wedge \text{intersects}(T_1, T_2) \wedge \\ \rightarrow (s_3, p_3, v_3)[\text{intersection}(T_1, T_2)]$$

Similarly, if the rule contains causality values C_1 and C_2 , an expression $\text{combine}(C_1, C_2)$ should be added to the rule's head to express the causality of the inferred statement. These operations can be quite complex, as there are multiple ways to express the same thing using qualifiers. For example, a temporal validity interval can be expressed, among other possibilities, with *start time* and *end time* or *start time* and *duration*. Furthermore, *start time* and *end time* may be left unspecified, typically representing $-\infty$ and $+\infty$, respectively. However, the complexity of dealing with the qualifier values is encapsulated in the data type operation definitions and does not appear in the rules and queries.

In [1] we propose a set of operations and their formal specification for a set of abstract qualifier types (context, sequence, causality, provenance, additional) that are close to some of the categories and subcategories of the present taxonomy.

6.3 Additional uses

In addition the taxonomy can help with

- Clarifying the description of each property when it is used as a qualifier. Many properties do not even have a textual description. Assigning a category to a qualifier is a good way to fix its use.
- Defining better user interfaces by organizing their elements according to the taxonomy (e.g. grouping the qualifiers that belong to the same category or organizing qualifiers in menus)

7 Related Works

In [12] the author distinguishes between additional and contextual qualifiers in Wikidata. He shows different solutions to represent and use these qualification in RDF. In particular, he

argues that contextualization qualifiers should not be shown to the users, but rather handled directly by tools such as reasoners according to a contextualization theory. By the way, our study tends to refute the assertion found in the article that “[...] many, perhaps most, uses of qualifiers in Wikidata and schema.org contextualize the underlying fact”. In fact, Figure 7 shows that contextualization represents less than 10% of the qualifications.

[7, 8] survey the notion of context in knowledge graphs in terms of implementation techniques (reification, higher-arity representation, and annotation) and operations on contexts. The survey refers to [16] that provides a uniform way to deal with all types of annotations by modeling annotation domains as commutative semi-rings with meet and join operations. Unlike our approach, this method does not distinguish between contextual (validity) annotations and other types of annotation; everything is treated uniformly. However, achieving this uniformity comes at the cost of a complex method for combining different domains, i.e. dealing with more than one annotation type in a query or reasoning task.

In [5] the authors present an adaptation of the `SchemaTree` recommender system for recommending qualifiers when creating new statements in the Wikidata knowledge base. The qualifier recommendation is not based on qualifier semantics or a pre-established taxonomy. Instead, it uses item and value type information, as well as co-occurring qualifiers.

It is also interesting to note that studies and guidelines on the design and construction of knowledge graphs generally do not address qualifier design, even in the context of property graphs. Similarly, studies of ontologies for or in Wikidata (e.g. [9] or [3]) are based on common ontology languages (RDFS, OWL, etc.) that do not support the notions of context, uncertainty, causality, any other type of qualification. Consequently they do not use qualifiers.

8 Conclusion

This work provides an in-depth quantitative and qualitative analysis of qualifier use in Wikidata. First, we show that the importance of a qualifier is determined by two independent factors: its frequency and the diversity of its usage. We measured this diversity using a variant of the diversity indexes employed in ecology. Our analysis of the 300 most important qualifiers resulted in the definition of a top-level taxonomy of qualifiers that reflects how qualifiers are actually used by the Wikidata contributors. In addition to the ‘contextual’ and ‘additional’ categories identified in previous studies we found that qualifiers can also convey epistemic information and participate in defining data structures. We also found that the previously identified categories could be refined into subcategories.

We evaluated the taxonomy in terms of coverage, showing how the most important qualifiers are well distributed across its categories. In terms of usefulness, we demonstrated how this taxonomy can help to solve two problems: of 1) selecting the right qualifiers when creating new statements; and 2) formulating correct queries and inference rules on the Wikidata graph.

From a more conceptual perspective, the taxonomy highlights the main modeling dimensions that should be considered when creating new knowledge graphs or when developing knowledge graph design methodologies.

Data Availability

The data (frequency dictionaries), diversity values, and classification results are available on GitHub at <https://github.com/cui-ke/wikidata-qualifiers>.

References

- [1] Sahar Aljalbout and Gilles Falquet. A Semantic Model for Historical Manuscripts. *arXiv preprint arXiv:1802.00295*, 2018.
- [2] Sahar Aljalbout, Gilles Falquet, and Didier Buchs. Handling wikidata qualifiers in reasoning. *arXiv preprint arXiv:2304.03375*, 2023.

- [3] Miquel Centelles and Núria Ferrer. Taxonomies and ontologies in wikipedia and wikidata: An in-depth examination of knowledge organization systems. *Hipertext.net*, pages 33–48, 05 2024.
- [4] Richard Cyganiak, David Wood, and Markus Lanthaler. Rdf 1.1 concepts and abstract syntax. Technical report, W3C, feb 2014.
- [5] Andrei Mihai Ducu and Michael Cochez. Qualifier Recommendation for Wikidata. In *In L.-A. Kaffee, S. Razniewski, K. Alghamdi, & H. Arnaout (Eds.), Wikidata 2023 The 4th Wikidata Workshop: Proceedings of the Wikidata Workshop 2023 co-located with 22nd International Semantic Web Conference (ISWC 2023) Athens, Greece, November 13, 2023*, volume 3640 of *CEUR Workshop Proceedings*, pages 1–12, Athens, 2023. CEUR-WS.org.
- [6] José M Giménez-García, Antoine Zimmermann, and Pierre Maret. NdFluents: An Ontology for Annotated Statements with Inference Preservation. In *European Semantic Web Conference*, pages 638–654. Springer, 2017.
- [7] Aidan Hogan. Context in graphs. In *Proceedings of the 1st International Workshop on Conceptualized Knowledge Graphs. RWTH Aachen University, Aachen*, 2018.
- [8] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia D’Amato, Gerard de Melo, Claudio Gutierrez, José Emilio Labra Gayo, Sabrina Kirrane, Sebastian Neumaier, Axel Polleres, Roberto Navigli, Axel-Cyrille Ngonga Ngomo, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. Knowledge Graphs. 2020.
- [9] Lucie-Aimée Kaffee, Simon Razniewski, Pavlos Vougiouklis, Valentina Anita Carrero, Paul Groth, Valentina Presutti, Lucie Kaffee, Simon Razniewski, and Pavlos Vougiouklis. Empirical ontology design patterns and shapes from wikidata. *Semantic Web*, 15(6):2293–2317, 2024.
- [10] MediaWiki contributors. Wikibase/datamodel. <https://www.mediawiki.org/wiki/Wikibase/DataModel>, 2025. Accessed: 2025-08-14.
- [11] Vinh Nguyen, Olivier Bodenreider, and Amit Sheth. Don’t like RDF reification?: making statements about statements using singleton property. In *Proceedings of the 23rd international conference on World wide web*, pages 759–770. ACM, 2014.
- [12] Peter F Patel-Schneider. Contextualization via qualifiers. In *Contextualized Knowledge Graphs @ ISWC 2018*, 2018.
- [13] Peter F Patel-Schneider. Contextualization via Qualifiers. In *First International Workshop on Contextualized Knowledge Graphs at ISWC 2018*, 2018.
- [14] B. Porter, V. Lifschitz, and F. Van Harmelen, editors. *Handbook of knowledge representation (1st ed)*. Elsevier, 2008.
- [15] Chris Tofallis. Add or multiply? a tutorial on ranking and choosing with multiple criteria. *INFORMS. Transactions on Education*, 14(3):109–119, 2014.
- [16] Antoine Zimmermann, Nuno Lopes, Axel Polleres, and Umberto Straccia. A general framework for representing, reasoning and querying with annotated semantic web data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 11:72–95, 2012.