

# Psychiq and Wwyyzzerdd: Wikidata completion using Wikipedia

Daniel Erenrich

*Sunnyvale, CA, United States*

*E-mail: daniel@erenrich.net*

**Solicited reviews:** Fariz Darari; anonymous reviewer; Dimitris Kontokostas

**Abstract.** Despite its size, Wikidata remains incomplete and inaccurate in many areas. Hundreds of thousands of articles on English Wikipedia have zero or limited meaningful structure on Wikidata. Much work has been done in the literature to partially or fully automate the process of completing knowledge graphs, but little of it has been practically applied to Wikidata. This paper presents two interconnected practical approaches to speeding up the Wikidata completion task. The first is Wwyyzzerdd, a browser extension that allows users to quickly import statements from Wikipedia to Wikidata. Wwyyzzerdd has been used to make over 100 thousand edits to Wikidata. The second is Psychiq, a new model for predicting instance and subclass statements based on English Wikipedia articles. Psychiq’s performance and characteristics make it well suited to solving a variety of problems for the Wikidata community. One initial use is integrating the Psychiq model into the Wwyyzzerdd browser extension.

**Keywords:** Wikidata, Wikipedia, Browser extension, Knowledge graph completion

## 1. Introduction

Wikidata is a collaborative knowledge base that is widely used as a source of structured data [1]. It is a multi-lingual, open-source database that is maintained by the Wikimedia Foundation and is closely linked to Wikipedia. Wikidata stores structured data about entities and links them to the Wikipedia versions in various languages about those entities.

Wikidata’s ontology is community moderated and is primarily defined by two kinds of elements: items and properties. Wikidata items are abstract entities which represent conceptual or material entities. This includes people, places and ideas. Properties are relations between entities or other pieces of data. This includes relations, such as birthdate, father, place of birth or height. Properties and items are assigned codes of the form P## or Q## respectively.

Probably the most important properties in the Wikidata ontology are “instance of” (denoted P31) and “subclass of” (P279). These statements form the backbone of the Wikidata ontology and, as others have noted, are effectively equivalent to `rdf:type` and `rdfs:subClassOf` [2]. P31 is used to express ideas, such as “Obama is a human” or “Rap is a music genre”. P279 is used to express ideas, such as “businesses are a kind of organization” or “film is a kind of artwork”.

While Wikidata is a separate project from Wikipedia, a lot of cross pollination occurs between the two platforms and much of the new content on Wikidata is sourced from Wikipedia. Despite this, it is still the case that many high quality articles on Wikipedia have little or no information on Wikidata. As of November 2022, there are 359,581 items on Wikidata corresponding to English Wikipedia articles that have no P31 or P279 statements. Roughly




Fig. 1. The ratio of pages on Wikidata to active users has steadily increased, indicating that than the growth in articles has outpaced the growth in the active user base on Wikidata.

60 thousand of those items have no statements whatsoever and many more only have uninformative external IDs imported for them. This is an underestimate of the number of English Wikipedia articles that are not well represented in Wikidata because many articles do not even have a corresponding item. Because Wikipedia has generally stricter notability inclusion guidelines than Wikidata, these items Wikidata is lacking coverage of are especially important.

Quantifying the importance of the zero-context items can be approximated by considering the daily web traffic for English Wikipedia articles whose Wikidata items lack both P31 and P279 statements. On March 1, 2023 these main namespace articles together had over 6.1 million views<sup>1</sup> out of a total daily English Wikipedia traffic of 89.3 million.

Improving or automating the process of migrating information from Wikipedia to Wikidata would improve Wikidata greatly. Today manual human driven editing is already a minority of all work done on Wikidata. Nearly half of all edits to Wikidata are performed by bots<sup>2</sup> and a large fraction of the remaining human-driven edits are made using automated tooling. Despite this, millions of statements are likely missing from Wikidata [3]. While it has been found to be more up-to-date than other knowledge graphs [4], Wikidata lags behind English Wikipedia on many subjects.

Wikipedia is often able to absorb new facts or concepts before Wikidata does because of its larger and more active editor and user-base. English Wikipedia alone has over 3 times more active editors than Wikidata<sup>3</sup> and these statistics already overestimate the actual participation on Wikidata. This measure of “active editors” (five or more edits) greatly favors Wikidata’s editor count as some operations on Wikidata require many successive edits. Additionally, most Wikidata editors in June 2023 who used the Wikipedia or Wikidata user-interfaces only edited sitelink information. Sitelink information is useful in that it ties Wikidata items to Wikipedia articles but it does not contain any statements (the kind of structured information this paper is concerned with).

In a typical 24-hour period ending January 2, 2023, 967,292 edits were made to Wikidata but only 85,779 (8.9%) of those edits were made using one of the Wikidata web interfaces<sup>4</sup>. The rest were made either by bots or through

<sup>1</sup>Using the pageview methodology described in <https://dumps.wikimedia.org/other/pageviews/readme.html>

<sup>2</sup>[https://stats.wikimedia.org/#/wikidata.org/contributing/edits/normalbar|2-year|editor\\_type~anonymous\\*group-bot\\*name-bot\\*user|monthly](https://stats.wikimedia.org/#/wikidata.org/contributing/edits/normalbar|2-year|editor_type~anonymous*group-bot*name-bot*user|monthly)

<sup>3</sup><https://stats.wikimedia.org>

<sup>4</sup><https://wikidata.wikiscan.org/hours/12/stats>

various other tools. Most edits are either fully automated or semi-automated. Semi-automated edits are generally batch changes made by importing information or performing some clerical task. 526,297 (54%) edits in the 24-hour period were made using QuickStatements which is a tool for bulk editing.

To appreciably increase the speed of Wikidata progress even more (semi)automated tools are needed. Figure 1 shows that the ratio of Wikidata items to active users is increasing, indicating that the community's size is not keeping up with the breadth of the Wikidata corpus<sup>5</sup>. While it is easy to just create more pages, curating a collection of this size is difficult. For a limited number of humans to curate a corpus of 100 million items better tooling is needed to more efficiently use their limited time.

The tools need to meet a number of requirements to be useful to the Wikimedia community. The most important of these requirements are: low latency, broad support for the most common Wikipedia patterns, flexibility to work with new Wikipedia content and high throughput of user actions. Low latency is critical because the tools are intended to be used in an online context and users will not tolerate a model that takes minutes to predict statements given an article. Broad support for Wikipedia patterns is important because the tool need to provide useful guidance in the most common cases. Flexibility is important because Wikipedia is constantly evolving and assumptions about, for example, the strict pattern of names of categories cannot be relied upon. For the tool to be high throughput it needs to require a minimal number of user interactions to perform common tasks.

This paper discusses the creation of two related tools to improve the process of importing data from Wikipedia into Wikidata: Wwyzzerdd and Psychiq. Wwyzzerdd is a user-interface browser extension that allows users to annotate Wikipedia links with properties and thus add new statements to Wikidata. Psychiq is a machine learning model that is integrated into the Wwyzzerdd UI that suggests, on the basis of Wikipedia's content, new P31 or P279 statements to be added to Wikidata.

## 2. Related work

Tools in this space fall somewhere on an spectrum between exclusively human directed and fully automated. It has previously been noted that a combination of automated and human-directed edits together has a strong positive impact on item quality in Wikidata [5].

Many of these tools in the Wikidata ecosystem have no mentions in the published literature even though some are heavily relied upon.

### 2.1. Automated extraction

For fifteen years DBpedia has largely automated the process of structuring Wikipedia using semantic web and linked data technologies [6]. DBpedia primarily does this by having volunteers contribute to the "DBpedia Mappings Wiki"<sup>6</sup> which provides a mapping from fields on Wikipedia templates to ontology classes. While DBpedia's approach is effective for many Wikipedia articles it is fairly reliant on articles that use infoboxes. Many articles, especially less mature ones, do not have infoboxes for DBpedia to harvest.

Others have attempted bootstrapping fully automated extraction of semantic information from Wikipedia. KYLIN is one such effort [7]. KYLIN is similar to the work presented here on Psychiq in that it learns a classifier from article content to classes. This work differs from KYLIN in that Psychiq relies heavily on the Wikipedia category system while the paper for KYLIN describes the Wikipedia categories as, "so flat and quirky that utility is low". KYLIN also was only trained on four classes while Psychiq targets 1000.

A much more ambitious effort is the Never Ending Language Learning (NELL) project [8]. The project aims to extract semantic meaning from the general web over an expanding number of learned relations and entities. The project appears to now be dormant and required training over 2,500 distinct learning tasks. NELL would be immensely overpowered for the use case considered here and may not meet Psychiq's latency requirements.

---

<sup>5</sup><https://stats.wikimedia.org>

<sup>6</sup>[http://mappings.dbpedia.org/index.php/Main\\_Page](http://mappings.dbpedia.org/index.php/Main_Page)

1 The Cat2Ax [9] system, like Psychiq, attempts to extract structured data from the category system that Wikipedia 1  
2 uses. Wikipedia has an extensive category system which is a hierarchy of classes that represent concepts, such as 2  
3 “Mountain ranges of Arizona”. Cat2Ax differs from the work done for Psychiq as it targets the DBpedia ontology 3  
4 instead of the Wikidata ontology. Cat2Ax also attempts to infer both type axioms and relation axioms (in the parlance 4  
5 of Wikidata that is equivalent to targeting more than just P31 and P279 statements). It also does not leverage the 5  
6 articles’ title information which is often critical for handling ambiguous cases and for which Cat2Ax employs a 6  
7 post-filtering step. One advantage Cat2Ax has is that in principle it is not limited to predicting high support types 7  
8 but for this paper’s goal of prediction of types in the Wikidata ontology this is not particularly important. Cat2Ax 8  
9 and subsequent work were used to create a new knowledge graph rather than feeding into Wikidata [10]. 9

10 Another automated attempt to populate statements is the Wikidata “NoclaimsBot”. This community maintained 10  
11 bot looks at the templates used in a Wikipedia article and, based on heuristics created by editors, infers statements 11  
12 for the item. For example, if the template “Infobox river” is used then the bot infers that the item for that page is an 12  
13 instance of a river (Q4022). This bot’s logic is simple and does not have mechanisms for resolving conflicts between 13  
14 multiple templates. The existence of this bot, however, impacts the quality of the training data in Wikidata. It means 14  
15 that much of the existing labels are being set using rigid rules. While it could, the bot does not set any references 15  
16 to indicate that the statements were inferred from templates used on Wikipedia. This means that there is no way to 16  
17 filter out automated statement additions from the training corpus. Since its inception in 2016, the NoclaimsBot has 17  
18 made 189,000 edits<sup>7</sup>. 18

19 In the same vein as the “NoclaimsBot” is the “Pi bot” [11]. While this bot does many kinds of tasks the relevant 19  
20 task for this paper is creating Wikidata items for new Wikipedia articles. Instead of being configured by the com- 20  
21 munity, it is controlled by a series of heuristics encoded in python code. It uses logic such as “if the title includes 21  
22 ’list of’ then the article is a Wikimedia list article (Q13406463)”. It also has more advanced heuristics where it 22  
23 attempts to deduce if the page is a biography and if it is it attempts to populate attributes like birth date, death date, 23  
24 gender and occupation using even more complex heuristics. “Pi bot” has made more than 8 million edits though not 24  
25 all of those are using these heuristics. 25

26 Kian [12] is similar to Psychiq in that it is a neural network designed to predict Wikidata statements. It was written 26  
27 in 2015 and does not leverage pre-trained language models. Instead it relies on heuristic features over categories 27  
28 and trains a binary classifier for a given statement type. Kian was used to add 100,000 statements to Wikidata [13]. 28  
29 Kian was gamified into a human-in-the-loop labeling application using the Wikidata Distributed Game framework<sup>8</sup> 29  
30 to handle cases where the model was not sufficiently confident. This game, however, is no longer online and the tool 30  
31 is not maintained. 31

32 While there is a lot of opportunity for automated extraction to improve Wikidata there isn’t consensus within 32  
33 the community about how to do so. Wikidata has thus far chosen not to import statement en masse from DBpedia 33  
34 even though doing so would be relatively easy. Efforts to add AI/ML inferred statements have been met with some 34  
35 skepticism over concerns over quality while simple heuristic solution have been preferred. There is a fear in the 35  
36 community of adding a large number of false statements which the human reviewers would never be able to audit. 36  
37 It is an open question how accurate a model would need to be to add newly inferred statements fully autonomously. 37  
38 38

## 39 2.2. Human-guided extraction 39

40 40  
41 Sztakipedia [14] is a browser extension that utilizes DBpedia Spotlight [15] to suggest improvements to 41  
42 Wikipedia. While these suggestions don’t explicitly “structure” the data or edit Wikidata the kinds of suggestions 42  
43 made are similar. Its design of modifying the user-interface of Wikipedia is also similar to Wwyyzzerdd. 43  
44 44

45 QuickStatements is community maintained tool that enable users to generate an offline batch of edits to Wikidata 45  
46 and execute them as a group via a web interface. While it has high usage numbers it serves as a convenient replace- 46  
47 ment for the Wikidata Python-API and does not aid the actual information extraction process. Users are left to do 47  
48 the data processing using some other tool. 48  
49 49

---

50 <sup>7</sup><https://xtools.wmflabs.org/ec/wikidata.org/NoclaimsBot> 50

51 <sup>8</sup><https://wikidata-game.toolforge.org/distributed> 51

PetScan is another community maintained tool that allows users to create filters for pages on a Wikipedia and then add statements on all matching items using QuickStatements. For example, a user could create a filter for all pages under the category “2018 films” which also use the template “Infobox film” and add the statement “instance of film” (P31 Q11424) to all the items. These batch workflows are human curated but it’s possible that without appropriate supervision incorrectly added statements would get through and cause bad category to statement associations to be inferred by Psychiq.

IntKB [16] is a human-in-the-loop system that uses natural language processing to suggest statements to add to Wikidata based on the text of Wikipedia. IntKB also has a companion Wikidata gadget which augments the user-interface to allow for quick editing of Wikidata<sup>9</sup>. While the tool is effective at extracting certain types of relations, it is particularly poorly suited for the important kinds of relations Psychiq is designed to extract. This is because instance/subclass relations are often not explicitly addressed in the opening text of an article (e.g. articles don’t call people “humans”). Also, while the tool is designed for human-in-the-loop work its community adoption is low with fewer than 1000 edits being made using the tool since its inception in September 2021. Of a random sample of 3657 of IntKB’s suggested statements, it suggested statements using P31 or P279 zero times.

The “Wikidata browser extension”<sup>10</sup> is the most similar tool to Wwyyzzerdd. However its focus is on extracting information from non-Wikipedia websites. When using the extension on Wikipedia, it has a few issues including not resolving redirects, not suggesting properties to link and not indicating which items are already linked. It has a significant amount of usage (roughly 70% of the daily edit volume of Wwyyzzerdd).

### 2.3. Gaps in related work

Of the existing automated extraction tools none meet all of the desired requirements. DBpedia is incompatible with Wikidata’s ontology. KYLIN supports only support four classes and so doesn’t support the breadth of Wikipedia’s content. Cat2Ax neither meets the latency or flexibility requirements as it is intended for offline batch analysis of Wikipedia dumps. Cat2Ax also only leverages the category information which reduces its accuracy. Kian is unmaintained and only works for a small number of classes. “NoclaimsBot” is inflexible and based on strict rules for category and template naming. “Pi bot” is entirely heuristic and supports a limited number of types.

Of the human-guided tools most are unsupported or solve a different problem. Sztakipedia no longer seems to be available to install. QuickStatements and PetScan don’t solve the problem of extracting structured data from a single Wikipedia article. IntKB almost never suggests P31 or P279 statements. The “Wikidata browser extension” isn’t optimized for extracting data from Wikipedia and so has higher latency and requires many more clicks to achieve the same thing as Wwyyzzerdd.

Given this, there is a need for new solutions to the problem.

## 3. Implementation

### 3.1. Wwyyzzerdd

Wwyyzzerdd is a browser extension for Firefox and Chrome that modifies the user-interface of Wikipedia. It is designed to make transcribing data from Wikipedia to Wikidata as efficient as possible. The main idea behind Wwyyzzerdd is that Wikipedia contains rich data about items that are not already present in Wikidata. If Wwyyzzerdd can make the process of structuring the data take just a few clicks then the coverage of Wikidata could be improved.

For example, a book might be properly tagged as a literary work but its authors and publishers may not be tagged in Wikidata. If the author is notable enough to have an English Wikipedia article, their name in the book’s article will almost always be a link to the author’s article. Wwyyzzerdd exploits the existence of this link.

---

<sup>9</sup><https://github.com/WikidataComplete/Wikidata-Complete-Gadget>

<sup>10</sup><https://github.com/fuddl/wd>

## InCryptid

An *unassessed* article from Wikipedia, the free encyclopedia  
 Wikidata: [InCryptid \(Q111943169\)](#), series of urban fantasy novels by American author Seanan McGuire  
 Aliases: None

**InCryptid** is a series of [urban fantasy](#) novels by American author [Seanan McGuire](#), published by [DAW Books](#).

Fig. 2. The Wwyyzzerdd user-interface showing linked entities (e.g. genre) in green and unlinked entities (e.g. publisher and author) in gray.

## InCryptid

An *unassessed* article from Wikipedia, the free encyclopedia  
 Wikidata: [InCryptid \(Q111943169\)](#), series of urban fantasy novels by American author Seanan McGuire  
 Aliases: None

**InCryptid** is a series of [urban fantasy](#) novels by American author [Seanan McGuire](#)

Contents [hide]	
1	<a href="#">Books</a>
2	<a href="#">Short stories</a>
3	<a href="#">Reception</a>
3.1	<a href="#">Awards and recognition</a>
4	<a href="#">Publication history</a>
5	<a href="#">References</a>

**Seanan McGuire**  
[Q3476794](#) · American writer

+ Property

Fig. 3. The Wwyyzzerdd user-interface shows a pop-up which allows the user to add a statement connecting the article with the target through an arbitrary Wikidata property.

Another common scenario is for the “NoClaimsBot” to populate a film’s value for P31 but not to import key people related to the works. Generally the director, producer, screenwriter and cast members of a film or TV series are notable enough to have their own pages on Wikipedia. So users could use Wwyyzzerdd to link those people to the item using the properties: P161 (cast member), P57 (director), P58 (screenwriter) or P162 (producer).

Wwyyzzerdd was implemented in Typescript [17] with the Parcel bundler<sup>11</sup> and using the yarn package manager<sup>12</sup>. The user-interface was built using React<sup>13</sup> with the Material UI component library<sup>14</sup>.

Using the Wikidata and Wikipedia APIs, Wwyyzzerdd interrogates all internal links in the current article and looks up the corresponding Wikidata items for them. It then renders an “orb” indicating if that item is already connected to the current item. If there is a property connecting the Wikidata item for the article to the item linked, the corresponding orb is green. Otherwise, the orb is gray. Figure 2 shows examples of both of these. Hovering over a green orb shows the list of properties linking the page to that item. If a user sees an entity which should be connected to the page’s item they can click on the orb and select the relevant property. Figure 3 shows an example of this. Properties are either suggested to the user using Wikibase’s integrated property suggestion algorithm (which suggests properties likely to be added the current item based on what’s currently in Wikidata) or the user can search for a desired by property by typing in the search box. For repetitive workflows (e.g. repeatedly tagging different book items with the same author), Wwyyzzerdd remembers the last property an item was attached with and preferentially suggests it.

By relying on the already populated links, Wwyyzzerdd avoids the well-known problem of entity-linkage between the Wikipedia text and Wikidata. This can sometimes be problematic as Wikipedia editors will sometimes not link common terms, such as the “United States”. Adding support for entity linkage techniques will likely be added later.

<sup>11</sup><https://parceljs.org/>

<sup>12</sup><https://yarnpkg.com/>

<sup>13</sup><https://react.dev/>

<sup>14</sup><https://mui.com/>

1 It is important that Wwyyzzerdd shows contextual information, such as the description and label of the linked 1  
2 item. This is because it is very common for the text of the link and the target of the link to not match. For example, 2  
3 Wikipedia articles often will say something similar to "Susan is an American author" where American is a hyperlink. 3  
4 Sometimes that hyperlink goes to the page for the "United State of America" and other times it will go to the article 4  
5 for the "American" ethnic group. Which value is linked determines the correct property to use. 5

6 Wwyyzzerdd also checks all external links in the Wikipedia article. External identifier properties in Wikidata are 6  
7 tagged with URL match patterns (P8966). These patterns are regular expressions which match URLs for the identi- 7  
8 fiers source. So for example, `orcid.org/(0000-{4}-{4}-{3}.)` matches ORCID iD URLs and extracts the 8  
9 underlying identifier. Wwyyzzerdd attempts to match all the Wikipedia external links on the URL match patterns. 9  
10 If a match is found then an orb is shown offering to let these user attach the identifier to the Wikidata item using the 10  
11 corresponding property. 11

12 To use the extension the user clicks on one of the orbs as shown in Fig 3. The extension uses the Wikidata API 12  
13 to suggest reasonable properties or the user can type to search. Once the user decides on a property and hits the 13  
14 plus button, Wwyyzzerdd makes two edits to Wikidata. First, it links the page's item to the target item through 14  
15 the selected property. Then it adds a reference to that statement indicating that the information was inferred by 15  
16 examining the current version of the Wikipedia article. 16

17 Wwyyzzerdd does not currently support adding qualifiers to statements but this is mostly needed for specific 17  
18 properties, such as "relative" (which should be qualified with "kinship to subject"). This may be solved in the future 18  
19 by adding custom support for relationships that require qualifiers (e.g. "uncle") which would automatically add the 19  
20 correct statements and qualifiers. 20  
21

### 22 3.2. *Psychiq* 22 23

24 *Psychiq*'s focus is exclusively on populating missing instance of (P31) and subclass of (P279) statements. This 24  
25 limiting decision was made for a number of reasons. First, as mentioned previously, P31 and P279 are the backbone 25  
26 of the Wikidata ontology and many other properties have constraints that require these values be populated. 26

27 Second, populating P31 and P279 statements is particularly difficult to do using Wwyyzzerdd alone. This is 27  
28 because it is rare for articles to link to items for their containing classes. For example, articles about people almost 28  
29 never link to the item for human (Q5) which is the correct Wikidata class. Many common classes for entities in 29  
30 Wikidata do not even have corresponding English Wikipedia articles. For example popular music songs are generally 30  
31 coded as Q105543609 but there is no corresponding English Wikipedia article for this concept. 31

32 Third, once the values of P31 or P279 are known Wikidata's property suggestion mechanism becomes much 32  
33 more capable. If Wikidata knows that the item is about a human then it intelligently suggests the properties gender, 33  
34 occupation, country of citizenship and date of birth. If instead the item is about a business it suggests headquarters 34  
35 location, country, inception and website. The property suggester used on Wikidata only has this level of intelligence 35  
36 when P31 or P279 are set.<sup>15</sup> 36

37 To simplify the machine learning process *Psychiq* only considers the title of an article and its categories. English 37  
38 Wikipedia has an extensive category system which alone allows for strong predictive performance. It is important to 38  
39 consider the title because some articles cannot be correctly classified only using categories. For example "General 39  
40 aviation in the United Kingdom" is impossible to classify as "aspect in a geographic region" (Q74817647) by 40  
41 only considering its categories of "Aviation in the United Kingdom" and "General aviation". In the future the first 41  
42 sentence of the article, the full-text of the article or category hierarchy information may also be added as a feature 42  
43 to the model. For now this information is excluded because of a desire to keep the inference latency and cost low. 43  
44 The full-text also cannot be included because the Distilbert model used in *Psychiq* does not support context lengths 44  
45 that can span the length of an entire Wikipedia article. 45

46 As was done for Cat2Ax, *Psychiq* filters the set of categories that are considered when making predictions. 46  
47 Categories were filtered out if they contained strings: 47  
48

- 49 – Short description 49

---

51 <sup>15</sup><https://github.com/wikimedia/mediawiki-extensions-PropertySuggester> 51

1	– Articles with	1
2	– All stub articles	2
3	– Wikidata	3
4	– Noindexed pages	4
5	– Redirects	5
6	– All articles	6
7	– dates	7
8	– Wikipedia articles	8
9	– wayback links	9
10	– Pages containing	10
11	– Articles containing	11
12	– Articles using	12
13	– Articles needing	13
14		14

These are largely administrative categories that do not contain substantial semantic information. In contrast with Cat2Ax though, categories containing the string "stub" are preserved as they contain important topical context. For example, the category “Swedish handball biography stubs” indicates that the article is likely about a person.

A primary goal of Psychiq is that it be fast enough to use in interactive contexts. While it’s unlikely that the model will be small enough to embed and run directly in the Wwyyzzerdd extension it must be small enough to host cheaply. To that end a small and fast model was selected: Distilbert-base-uncased.

Distilbert-base-uncased [18] was selected because of its size, latency and capacity. Distilbert is small given its performance, just a couple hundred MBs of weights, and was trained over the Wikipedia English corpus which makes it a good fit for this application. It is pared down version of the BERT model [19] which, while no longer state-of-the-art, has become a standard NLP baseline. Distilbert and BERT both leverage transformers models. Distilbert was trained by “distilling” BERT to a similarly shaped model with fewer layers.

In practice the model is hosted on huggingface.com and the model runs on an average input in 35ms once the instance is spun-up. Distilbert is trained exclusively on English so fine-tuning training/evaluation data for Psychiq is restricted to English Wikipedia.

Psychiq is restricted to predicting amongst the top 1000 most frequent P31 or P279 statements. Of all the English Wikipedia articles with P31 or P279 statements on the linked Wikidata items 93% of them use one of these top 1000 most frequent items. The one-thousandth most frequent statement occurred 356 times. Many of the items that do not use one of the 1000 most common statements would still be well covered by a top-1000 statements. For example, instances of “village of Japan” (Q4174776) would also be well represented by “village” (Q532).

The Hugging Face Transformers library [20] was used to train a sequence classification model for 1001 classes. This model feeds the text corresponding to a Wikipedia page through a BERT-like neural-network to produce an embedding of the text. This embedding is then fed into a final dense linear layer with a number of neurons equal to the number of classes and with a cross-entropy loss function.

As input to the model, we generate a text document for each Wikipedia page composed of a listing of all the categories separated by newlines followed by the title of the article. The order of the categories is left as-is in the article because Wikipedia assigns some small meaning to the ordering of the categories<sup>16</sup>. A sample document for the article “Bandaje Falls” would be:

```
Waterfalls of Karnataka
Tourist attractions in Dakshina Kannada district
Geography of Dakshina Kannada district
Bandaje Falls
```

<sup>16</sup>[https://en.wikipedia.org/wiki/Wikipedia: Categorization#General\\_conventions](https://en.wikipedia.org/wiki/Wikipedia: Categorization#General_conventions)






Fig. 4. Psychiq’s accuracy is reliably above 80%, especially in the top-5 accuracy context where its used in Wwyyzzerdd. The chart compares Psychiq’s performance against suggesting the most frequent statements (i.e. guessing).

Table 1  
Psychiq test set performance metrics

Metric	Value
Accuracy	84.6%
Precision (Weighted)	83.2%
Recall (Weighted)	84.6%
ROC AUC (Weighted, one-vs-rest)	0.996

The corpus data processing was done using BigQuery. Freely available dumps of Wikidata and English Wikipedia were loaded into Google Cloud Storage and then imported into BigQuery. In particular, the corpus of “enwiki-categorylinks”<sup>17</sup> was used to map English Wikipedia articles to categories. The “wikidata-nt-all”<sup>18</sup> corpus was downloaded and split into two. One dataset was a mapping of Wikidata items to statements and one was a mapping from Wikidata items to English Wikipedia articles.

The 1000 most common values for the property of P31 or P279 for items with corresponding English Wikipedia articles were computed and each was assigned a unique label value. A single “example” was generated for each English Wikipedia article, excluding redirects, with a specific value for P31 or P279. These examples were joined with the labels with a special 1001 label value introduced for statements outside of the top 1000.

The training and testing corpus of these documents and the accompanying labels is available on Hugging Face<sup>19</sup>.

The model is then asked to predict one of the top 1000 statements or none-of-the-above based on these documents. The correct prediction for that sample document would be “waterfall” (Q34038). Training was performed on Google Collab for a single epoch, due to budgetary constraints. Over 5.6 million examples were in the training set leaving another 10% out as a test set. The hyper-parameters for learning were a learning rate of 2e-5, a batch size of 32 and a weight decay of 0.01. The trained model is available for download on Hugging Face<sup>20</sup>. Scikit-learn was used to compute performance metrics [21]. Source code for downloading/preparing data and training the model is on GitHub<sup>21</sup>

<sup>17</sup><https://dumps.wikimedia.org/enwiki/latest/enwiki-latest-categorylinks.sql.gz>

<sup>18</sup><https://dumps.wikimedia.org/wikidatawiki/entities/latest-all.nt.bz2>

<sup>19</sup><https://huggingface.co/datasets/derenrich/psychiq2-dataset>

<sup>20</sup><https://huggingface.co/derenrich/psychiq2>

<sup>21</sup><https://github.com/derenrich/psychiq>

Table 2  
Most confused pairs of statements

Statement 1	Statement 2	Error frequency
instance of disambiguation page (Q4167410)	instance of human name disambiguation page (Q22808320)	1131
instance of written work (Q47461344)	instance of literary work (Q7725634)	1087
instance of gene (Q7187)	subclass of protein-coding gene (Q20747295)	1031
instance of musical work/composition (Q105543609)	instance of single (Q134556)	1087
unknown	instance of business (Q4830453)	976
instance of civil parish (Q1115575)	instance of village (Q532)	941

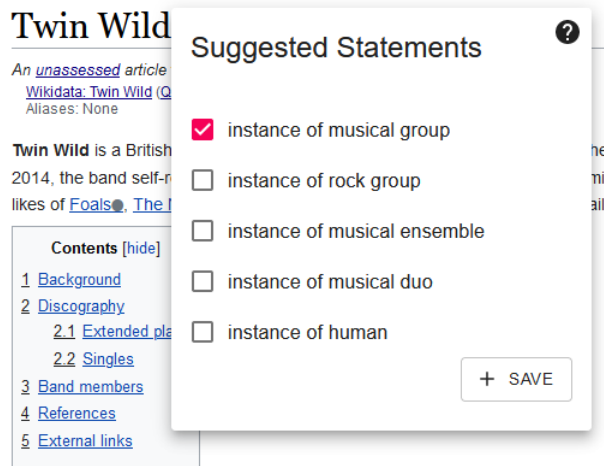


Fig. 5. Psychiq is integrated into Wwyyzzerdd as a drop-down that lets the user select statements to add to Wikidata.

Table 1 shows the key test set performance indicators. Qualitatively the performance is good enough for its intended human-in-the-loop use. This is especially true when considering the top-5 scenario where it is used in Wwyyzzerdd (users are shown the 5-top predictions from the model). Figure 4 shows the model’s accuracy in various top-K scenarios. The “guessing” baseline is the accuracy of always selecting the top-K most common labels. The confusion matrix indicates which pairs of statements the model is worst at distinguishing. The most confused pairs are presented in Table 2. Most of the pairs are self-explanatory and are the result of either very closely coupled concepts (a song and the single it was released on) or nested concepts (literary works are a kind of written work).

As noted by Shenoy [3], Wikidata often has trouble distinguishing between “subclass of” and “instance of” statements. This is corroborated by Psychiq’s confusion between instances of genes (DNA) and subclasses of protein-coding gene. An example of this is the gene the gene *DYNC1I1* (Q248215). While it is often sensible for an item to be both an instance and subclass of other items (e.g. rap is both a subclass of music and an instance of music genre) it does not make sense for an item to be an instance and subclass of the same thing (in this case kinds of genes). Areas where the ontology could be improved can be found by examining Psychiq’s errors.

Another common issue occurs when Wikidata and Wikipedia disagree. For example, for many of the Wikidata items about genes the corresponding English Wikipedia article is actually about the protein that the gene codes for. The Wikidata item for the gene *DYNC1I1* lists itself for the property “encoded by” (P702) which can’t be right because the item is a subclass of protein-coding gene and not a protein. When Wikidata and Wikipedia disagree what the topic is you cannot safely infer Wikidata statements from the Wikipedia article. This is a limitation of this approach.

As shown in Fig. 5 Psychiq was integrated into the Wwyyzzerdd user-interface by adding a drop-down next to the Wikipedia article’s title which shows the user the top few suggested statements for that article. The orb is green if one of the suggested statements is already present on the item.

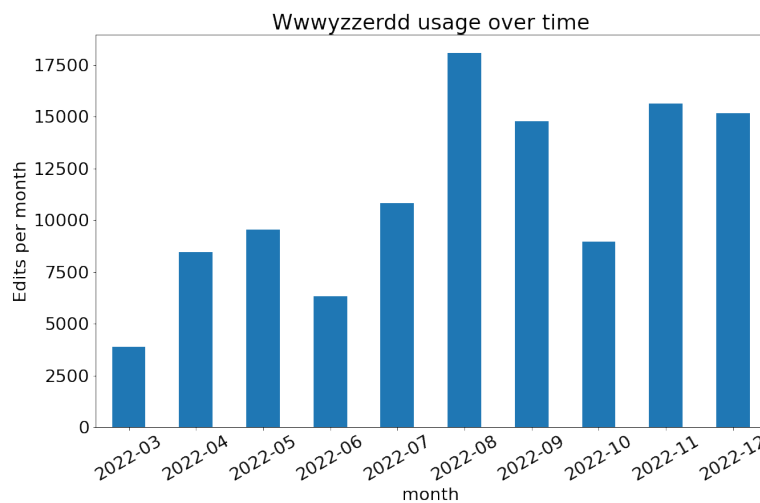


Fig. 6. Wwyyzerdd has seen strong and growing usage since logging began in March 2022.

## 4. Discussion

### 4.1. Wwyyzerdd

Wwyyzerdd has seen significant adoption. 111,655 edits were made using the tool in 2022 after edit logging began in March. Fig 6 shows a trend of increasing adoption over time. 69 distinct users tried the tool. The Chrome Web Store reports 58 installs<sup>22</sup> and the Mozilla Add-Ons website reports 25 users<sup>23</sup>.

While it is difficult to judge the quality or impact of the edits being made, one measure is the fraction of edits made that end up reverted. Good or productive edits tend not to be reverted. In total 0.69% of edits made using the tool were reverted. For comparison, on a typical day, January 1 2023, about 1.3% of all edits made using the Wikidata web UI were reverted. Investigating the reverted Wwyyzerdd edits individually suggests that a significant fraction are users reverting themselves after making an error. A possible explanation for the lower revert rate of Wwyyzerdd users is that the userbase of Wwyyzerdd is self selected to the most active and knowledgeable users.

Evaluation of the Wwyyzerdd edits using standard linked data quality measures[22] was not performed. The revert measure used above is an informal measure of quality and further study should be conducted to quantify the quality of the edits made using this tool.

While the tool was originally only available in English it works on all language editions of Wikipedia and has seen adoption across non-English Wikipedias. 43% of edits were made from the English Wikipedia, 36% were from the Japanese Wikipedia and, in descending order, most of the rest of the edits were from French, Belarusian, Russian, Chinese and Dutch Wikipedias.

Wwyyzerdd has been used to add statements for a diverse set of properties. The most common property added using Wwyyzerdd was P19 (place of birth) at 7.8% of all edits. This is followed by P31 (instance of) at 5.4% (this includes edits made using the Psychiq integration), P69 (educated at) at 5.3% and P161 (cast member) at 4.9%. 14.2% of statements added involved external identifiers and the most commonly added external identifier was P2002 (Twitter username).

A fundamental limitation of Wwyyzerdd and Psychiq is that they rely on the accuracy and reliability of (English) Wikipedia. The sources that Wwyyzerdd extracts from Wikipedia are not to the underlying sources used in Wikipedia but to Wikipedia itself. This reliance on secondary sources can be problematic especially if the data is to be reused within Wikipedia itself (creating a referential loop).

<sup>22</sup><https://chrome.google.com/webstore/detail/wwyyzerdd-for-wikidata/gfidggfngdnaalpihbjnfbkfiniookc>

<sup>23</sup><https://addons.mozilla.org/en-CA/firefox/addon/wwyyzerdd-for-wikidata/>

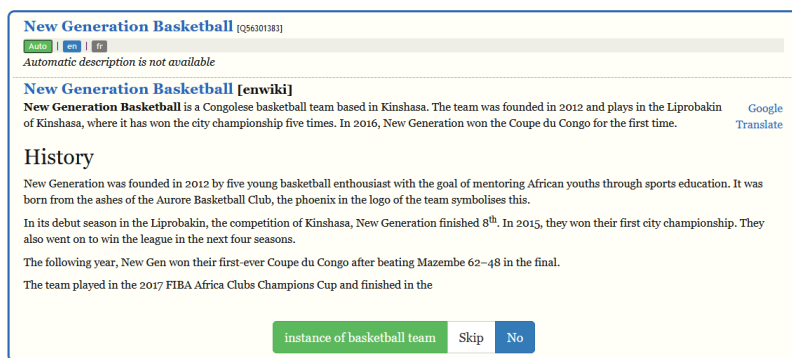


Fig. 7. Using the Wikidata Distributed Game framework, Psychiq’s top predictions are “gamified” so that users can easily validate and execute its suggestions.

Referencing Wikipedia is convenient to the editor and better than providing no reference but is not as good as it could be. A future version of the software could allow users to import primary sources from Wikipedia directly into Wikidata. For example, a future version could allow users to drag-and-drop a reference footnote onto an orb to export that reference directly to Wikidata.

Another limitation of Wwyyzzerdd is that subsequent updates to Wikipedia are not automatically imported to Wikidata. This means that errors corrected in one place are not corrected in the other. A fully automated text to structured data system would allow updates to flow immediately from one to the other, however it is not clear such a system would be accepted by the Wikidata community.

Without user studies under controlled conditions it is impossible to fully assess whether users like the experience of using Wwyyzzerdd and are made more productive by it. However, informal polling of the most frequent users suggests enthusiasm for the browser extension and a desire for it to take on more capabilities. The growth in monthly edit counts since launch also suggests that Wwyyzzerdd is successfully meeting a niche need in the Wikidata community.

In the future Wwyyzzerdd may be extended to work on non-Wikipedia websites. For that to happen a named entity extraction and linkage system would need to be integrated to allow users to select mentions of entities in general Web text (this isn’t needed on Wikipedia as text is already linked to “entities” through intra-wiki hyperlinks). Users would validate the linkage back to Wikidata and draw edges between entities and select properties to connect them. The URL of the current web page could then be used as the reference for the added statement.

#### 4.2. Psychiq

Comparing Psychiq to past scholarly work is difficult as most of the past work has been evaluated against DB-Pedia. Cat2Ax reports a 95.7% precision for the equivalent “type assertion” task. This is considerably higher than 83.2% reported in Table 1, however their performance metric is computed differently. Instead of checking for strict label equality Cat2Ax used human evaluators to judge the output as “correct or incorrect”. It is very common for Psychiq to assign a label that doesn’t strictly match what it is in Wikidata but that a human evaluator would judge as correct (indeed table 2 shows that the most common errors are of this type). So the precision of Psychiq is likely underestimated under Cat2Ax’s metric.

To get a better understanding of Psychiq’s model performance a separate experiment was performed. In addition to being integrated into Wwyyzzerdd, Psychiq was converted into a simple “game-like” web interface<sup>24</sup>. Figure 7 shows what the user interface for this game looks like. The tool shows users a Wikidata item, its corresponding English Wikipedia article and a suggested statement. The player can then accept the statement as correct, mark the statement as incorrect or skip the statement if they are unsure. The game was seeded with the 30 thousand proposed

<sup>24</sup><https://wikidata-game.toolforge.org/distributed/#game=84>

1 statements where Psychiq was most confident<sup>25</sup>. After 10,595 statements were evaluated by human volunteers, 2 over 97.7% of suggested statements were accepted. Initial human feedback on this tool suggested the model was 3 performing well.

4 The Psychiq model is robust to missing information. Newly created articles often lack categories and only have 5 titles so a system like Cat2Ax would not be applicable. As an example, for the article “University of Vechta” without 6 categories the model still accurately suggests that it may be a university. The model, however, will not work with 7 only category information because the model assumes the name of the category is the name of the article. For 8 example, just the category “Mosques in Jordan” is predicted to be an instance of “aspect in a geographic region” 9 (Q74817647) when it should have predicted “mosque” (Q32815). Fortunately, it is not possible for an article to 10 have categories but no title.

11 Another use case of Psychiq is identifying errors in Wikidata. By looking for instances where the model is very 12 confident and disagrees with what is currently stated in Wikidata, Psychiq can find probable errors. Unfortunately 13 doing this also produces many uninteresting conflicts where, for instance, the model identifies the item as a road 14 but it’s labeled as a highway. By filtering out uninformative pairs based on their frequency, one is left with more 15 actionable discrepancies. Of the remaining top fifty cases where Psychiq is most confident Wikidata is in error<sup>26</sup> 16 manual review found all but one to be actual errors in Wikidata. The most probable detected errors were caused 17 by English Wikipedia pages being converted from disambiguation pages to articles about specific people without 18 updating the corresponding Wikidata item.

19 Wwyyzzerdd already works in multiple languages, and it would not be difficult to also make Psychiq work 20 for multiple languages. All that would be required is expanding the training set to multiple language Wikipedias 21 and using a language model trained on a multi-lingual corpus. Using Wwyyzzerdd in a multi-lingual context does 22 introduce problems if the different linked Wikipedia articles are about subtly different subjects.

23 A major difficulty in Wikidata is establishing and communicating the desired ontology for various classes of 24 information. It is not obvious to new users that “Beat It” (Q210218) should be classed as an instance of “musical 25 work/composition” (Q105543609) and not the more natural sounding “song” (Q7366). Another common com- 26 mon confusion is between “bibliography” (Q1631107) and “bibliography” (Q134995) which are two different 27 concepts with the exact same name. Today this information is communicated through Wikiprojects which write up 28 documentation<sup>27</sup> or through model items<sup>28</sup> which are maintained examples of the desired way to model information. 29 Instead, with Psychiq, the correct way to model an item can be suggested using machine learning. If the community 30 wants to change how to represent something, they just needed to update all the existing items to the new ontology 31 and then the next time the model is retrained the suggestions will automatically reflect the new ontological model. 32

33 While Psychiq’s performance numbers look encouraging it is likely that selection bias makes the number appear 34 more positive than they may otherwise be. Psychiq is only trained on Wikidata items that have already been catego- 35 rized and items that are difficult to categorize are less likely to have been categorized. Thus the dataset that Psychiq 36 is trained on is unrepresentatively easy to classify.

37 Psychiq’s performance could likely be greatly expanded by using a more powerful and modern large-language 38 model. Additionally the model could be given access to the entire article’s text instead of just the title and categories. 39 For now the cost and latency of such an approach makes it prohibitive. 40

## 41 5. Conclusion 42

43 This paper has discussed two related approaches for improving the user-experience of Wikidata completion us- 44 ing the textual content of Wikipedia. This work represents a step beyond the current dominant tools in the Wikidata 45 46

---

47 <sup>25</sup>Given the limited human volunteer time it does not make sense to have people evaluate statements of lower confidence even though doing 48 so would provide a more comprehensive view on Psychiq’s accuracy

49 <sup>26</sup><https://www.wikidata.org/wiki/User:BrokenSegue/PsychiqConflicts>

50 <sup>27</sup>[https://www.wikidata.org/wiki/Wikidata:WikiProject\\_Music](https://www.wikidata.org/wiki/Wikidata:WikiProject_Music)

51 <sup>28</sup>[https://www.wikidata.org/wiki/Wikidata:Model\\_items](https://www.wikidata.org/wiki/Wikidata:Model_items)

ecosystem. The state of the art of tooling in the Wikidata automation space still remains behind the work of the competition. DBpedia’s approaches to automation fifteen years ago remain more advanced than many of the approaches being actively used in Wikidata. This is partly due to a difference in philosophy over how contributions should be made and partly a lack of investment. Hopefully this and similar parallel efforts will move Wikidata towards the state-of-the-art.

## References

- [1] D. Vrandečić and M. Krötzsch, Wikidata: a free collaborative knowledgebase, *Communications of the ACM* **57**(10) (2014), 78–85.
- [2] A. Ismayilov, D. Kontokostas, S. Auer, J. Lehmann, S. Hellmann et al., Wikidata through the Eyes of DBpedia, *Semantic Web* **9**(4) (2018), 493–503.
- [3] K. Shenoy, F. Ilievski, D. Garijo, D. Schwabe and P. Szekely, A Study of the Quality of Wikidata, *Journal of Web Semantics* **72** (2022), 100679.
- [4] S.G. Pillai, L.-K. Soon and S.-C. Haw, Comparing DBpedia, Wikidata, and YAGO for Web Information Retrieval, in: *Intelligent and Interactive Computing*, V. Piuri, V.E. Balas, S. Borah and S.S. Syed Ahmad, eds, Springer Singapore, Singapore, 2019, pp. 525–535. ISBN 978-981-13-6031-2.
- [5] A. Piscopo, C. Phethean and E. Simperl, What makes a good collaborative knowledge graph: group composition and quality in wikidata, in: *International Conference on Social Informatics*, Springer, 2017, pp. 305–322.
- [6] J. Lehmann, R. Isele, M. Jakob, A. Jentzsch, D. Kontokostas, P.N. Mendes, S. Hellmann, M. Morsey, P. Van Kleef, S. Auer et al., Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia, *Semantic web* **6**(2) (2015), 167–195.
- [7] F. Wu and D.S. Weld, Autonomously semantifying wikipedia, in: *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, 2007, pp. 41–50.
- [8] T. Mitchell, W. Cohen, E. Hruschka, P. Talukdar, J. Betteridge, A. Carlson, B. Dalvi, M. Gardner, B. Kisiel, J. Krishnamurthy, N. Lao, K. Mazaitis, T. Mohamed, N. Nakashole, E. Platanios, A. Ritter, M. Samadi, B. Settles, R. Wang, D. Wijaya, A. Gupta, X. Chen, A. Saparov, M. Greaves and J. Welling, Never-Ending Learning, in: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, 2015.
- [9] N. Heist and H. Paulheim, Uncovering the Semantics of Wikipedia Categories, in: *The Semantic Web – ISWC 2019*, C. Ghidini, O. Hartig, M. Maleshkova, V. Svátek, I. Cruz, A. Hogan, J. Song, M. Lefrançois and F. Gandon, eds, Springer International Publishing, Cham, 2019, pp. 219–236. ISBN 978-3-030-30793-6.
- [10] N. Heist and H. Paulheim, Information extraction from co-occurring similar entities, in: *Proceedings of the Web Conference 2021*, 2021, pp. 3999–4009.
- [11] M. Peel, Pi Bot. <https://github.com/mpeel/wikicode>.
- [12] A. Sarabadani, Kian. <https://github.com/Ladsgroup/Kian>.
- [13] Amir Sarabadani and Daniel Erenrich, Re: Kian, 2023.
- [14] M. Héder and P. Mendes, Round-trip semantics with Sztakipedia and DBpedia Spotlight (2012). doi:10.1145/2187980.2188048.
- [15] P.N. Mendes, M. Jakob, A. Garcia-Silva and C. Bizer, DBpedia Spotlight: Shedding Light on the Web of Documents, in: *Proceedings of the 7th International Conference on Semantic Systems (I-Semantics)*, 2011.
- [16] B. Kratzwald, G. Künpeng, S. Feuerriegel and D. Diefenbach, IntKB: A Verifiable Interactive Framework for Knowledge Base Completion, in: *Proceedings of the 28th International Conference on Computational Linguistics*, International Committee on Computational Linguistics, Barcelona, Spain (Online), 2020, pp. 5591–5603. doi:10.18653/v1/2020.coling-main.490. <https://aclanthology.org/2020.coling-main.490>.
- [17] G. Bierman, M. Abadi and M. Torgersen, Understanding typescript, in: *European Conference on Object-Oriented Programming*, Springer, 2014, pp. 257–281.
- [18] V. Sanh, L. Debut, J. Chaumond and T. Wolf, DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter, *ArXiv abs/1910.01108* (2019).
- [19] J. Devlin, M. Chang, K. Lee and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, J. Burstein, C. Doran and T. Solorio, eds, Association for Computational Linguistics, 2019, pp. 4171–4186. doi:10.18653/v1/n19-1423.
- [20] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, C. Ma, Y. Jernite, J. Plu, C. Xu, T. Le Scao, S. Gugger, M. Drame, Q. Lhoest and A.M. Rush, Transformers: State-of-the-Art Natural Language Processing, Association for Computational Linguistics, 2020, pp. 38–45. <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [21] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay, Scikit-learn: Machine Learning in Python, *Journal of Machine Learning Research* **12** (2011), 2825–2830.
- [22] A. Zaveri, A. Rula, A. Maurino, R. Pietrobon, J. Lehmann and S. Auer, Quality assessment for linked data: A survey, *Semantic Web* **7**(1) (2016), 63–93.