

Psychiq and Wwyyzzerdd: Wikidata completion using Wikipedia

Daniel Erenrich

Sunnyvale, CA, United States

E-mail: daniel@erenrich.net

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.

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 things like people, places and ideas. Properties are relations between entities or other pieces of data. This includes relations like birth date, 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 like “Obama is a human” or “Rap is a music genre”. P279 is used to express ideas like “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 60 thousand of those items have no statements whatsoever. 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

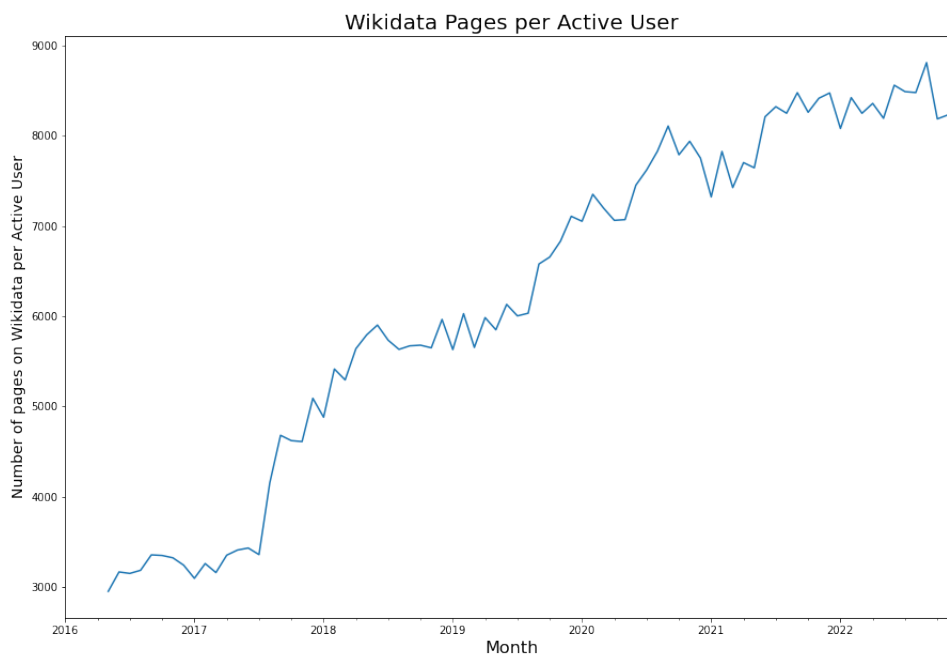


Fig. 1. The number of items on Wikidata has grown much faster than the number of active users on Wikidata.

item. Because Wikipedia has stricter notability inclusion guidelines than Wikidata, these items Wikidata is lacking coverage of are especially important.

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 bots¹ 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.

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 interfaces². The rest were made either by bots or through 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 growth in active users is not keeping up with the breadth of the Wikidata corpus³. 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.

This paper discusses the creation of two related tools to improve the process of importing data from Wikipedia into Wikidata: Wwyyzzerdd and Psychiq. Wwyyzzerdd 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 Wwyyzzerdd UI that suggests, on the basis of Wikipedia's content, new statements to be added to Wikidata.

¹https://stats.wikimedia.org/#/wikidata.org/contributing/edits/normalbar|2-year|editor_type~anonymous*group-bot*name-bot*user|monthly

²<https://wikidata.wikiscan.org/hours/12/stats>

³<https://stats.wikimedia.org>

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"⁴ 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.

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

Another automated attempt to populate statements is the Wikidata "NoclaimsBot". This community maintained bot looks at the templates used in a Wikipedia article and, based on heuristics created by editors, infers statements for the item. For example, if the template "Infobox river" is used then the bot infers that the item for that page is an instance of a river (Q4022). This bot's logic is simple and does not have mechanisms for resolving conflicts between multiple templates. The existence of this bot, however, impacts the quality of the training data in Wikidata. It means that much of the existing labels are being set using rigid rules. While it could, the bot does not set any references to indicate that the statements were inferred from templates used on Wikipedia. This means that there is no way to filter our automated statement additions from the training corpus. Since its inception in 2016, the NoclaimsBot has made 189,000 edits⁵.

In the same vein as the "NoclaimsBot" is the "Pi bot"[10]. While this bot does many kinds of tasks the relevant task for this paper is creating Wikidata items for new Wikipedia articles. Instead of being configured by the community, it is controlled by a series of heuristics encoded in python code. It uses logic like "if the title includes 'list of' then the article is a Wikimedia list article (Q13406463)". It also has more advanced heuristics where it attempts to deduce if the page is a biography and if it is it attempts to populate attributes like birth date, death date, gender and occupation using even more complex heuristics. "Pi bot" has made more than 8 million edits though not all of those are using these heuristics.

Kian [11] is similar to Psychiq in that it is a neural network designed to predict Wikidata statements. It was written in 2015 and does not leverage pre-trained language models. Instead it relies on heuristic features over categories and trains a binary classifier for a given statement type. Kian was used to add 100,000 statements to Wikidata[12].

⁴http://mappings.dbpedia.org/index.php/Main_Page

⁵<https://xtools.wmflabs.org/ec/wikidata.org/NoclaimsBot>

Kian was gamefied into a human-in-the-loop labeling application using the Wikidata Distributed Game framework⁶ to handle cases where the model was not sufficiently confident.

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

2.2. Human-guided extraction

Sztakipedia[13] is a browser extension that utilizes DBpedia Spotlight[14] to suggest improvements to Wikipedia. While these suggestions don't explicitly "structure" the data or edit Wikidata the kinds of suggestions made are similar. Its design of modifying the user-interface of Wikipedia is also similar to Wwyzzerdd.

PetScan is a 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 [15] 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 Wikidata⁷. 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 less than 1000 edits being made using the tool since its inception in September 2021.

The "Wikidata browser extension"⁸ is the most similar tool to Wwyzzerdd. 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 Wwyzzerdd).

3. Implementation

3.1. Wwyzzerdd

Wwyzzerdd is a manifest V2 compatible 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 Wwyzzerdd is that Wikipedia contains rich data about items that are not already present in Wikidata. If Wwyzzerdd 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. Wwyzzerdd exploits the existence of this link.

Using the Wikidata and Wikipedia APIs, Wwyzzerdd 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

⁶<https://wikidata-game.toolforge.org/distributed>

⁷<https://github.com/WikidataComplete/Wikidata-Complete-Gadget>

⁸[urlhttps://github.com/fuddl/wd](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 like genre and unlinked entities like publisher and author.

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	Books
2	Short stories
3	Reception
3.1	Awards and recognition
4	Publication history
5	References




Fig. 3. The Wwyyzzerdd user-interface allows you to link a book to its author.

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. 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.

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 like the “United States”. Adding support for entity linkage techniques will likely be added later.

It is important that Wwyyzzerdd shows contextual information like the description and label of the linked 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, Wikipedia articles often will say something like "Susan is an American author" where American is a hyperlink. Sometimes that hyperlink goes to the page for the United State of America and other times it will go to the article for the American ethnic group. Which value is linked determines the correct property to use.

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

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

Wwyyzzerdd does not currently support adding qualifiers to statements but this is mostly only needed for specific properties like “relative“ (which should be qualified with “kinship to subject”). This may be solved in the future by

adding custom support for relationships that require qualifiers like “uncle” which would automatically add the correct statements.

3.2. Psychiq

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

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

Third, once the values of P31 or P279 are known Wikidata’s property suggestion mechanism becomes much more capable. If Wikidata knows that the item is about a human then it intelligently suggests the properties gender, occupation, country of citizenship and date of birth. If instead the item is about a business it suggests headquarters location, country, inception and website. The property suggester used on Wikidata only has this level of intelligence encoded for P31 and P279⁹.

To simplify the machine learning process Psychiq only considers the title of an article and its categories. English Wikipedia has an extensive category system which alone allows for strong predictive performance. It is important to consider the title because articles like “General aviation in the United Kingdom” are impossible to classify as “aspect in a geographic region” (Q74817647) by only considering its categories of “Aviation in the United Kingdom” and “General aviation”. In the future the first sentence of the article or category hierarchy information may also be added as a feature to the model.

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

- Short description
- Articles with
- All stub articles
- Wikidata
- Noindexed pages
- Redirects
- All articles
- dates
- Wikipedia articles
- wayback links
- Pages containing
- Articles containing
- Articles using
- Articles needing

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 [16] was selected because of

⁹<https://github.com/wikimedia/mediawiki-extensions-PropertySuggester>

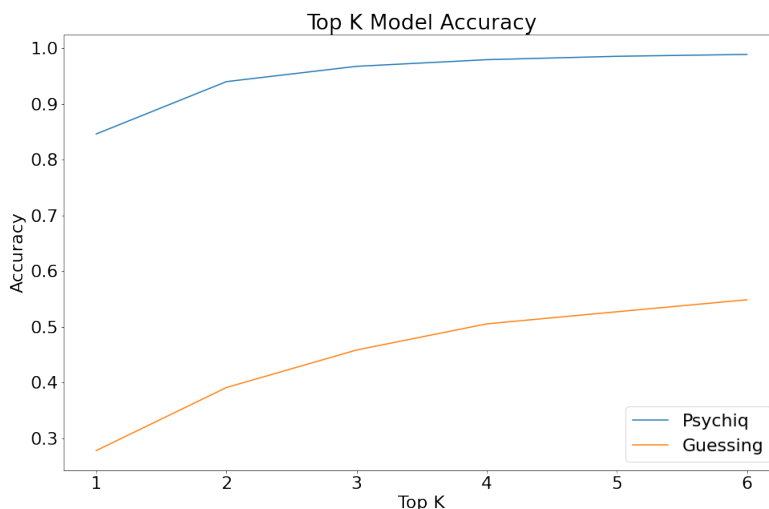


Fig. 4. The model’s accuracy is good but if you consider its top 5 predictions its performance is much better. The chart also shows the performance of just suggesting the most frequent statements.

its size and capacity. 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. This language model is trained exclusively on English so training/evaluation data 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 [17] was used to train a sequence classification model for 1001 classes. 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 categories¹⁰. A sample document for the article “Bandaje Falls” would look like:

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

The training and testing corpus of these documents and the accompanying labels is available on Hugging Face¹¹.

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 for a single epoch over 5.6 million examples 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 Face¹². Scikit-learn was used to compute performance metrics[18]. Source code for downloading/preparing data and training the model is on GitHub¹³

¹⁰https://en.wikipedia.org/wiki/Wikipedia:Category:General_conventions

¹¹<https://huggingface.co/datasets/derenrich/psychiq2-dataset>

¹²<https://huggingface.co/derenrich/psychiq2>

¹³<https://github.com/derenrich/psychiq>

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

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

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 Wwyzzerdd (users are shown the 5-top predictions from the model). Figure 4 shows the model’s accuracy in various top-K scenarios. 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 the Psychiq’s confusion between instances of genes (DNA) and subclasses of protein-coding gene. An example of this is the gene the gene *DYNC111* (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. For example, the Wikidata item for the gene *DYNC111* 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.

As shown in Fig. 5 Psychiq was integrated into the Wwyzzerdd 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.

4. Discussion

4.1. Wwyzzerdd

Wwyzzerdd 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 installs¹⁴ and the Mozilla Add-Ons website reports 25 users¹⁵.

¹⁴<https://chrome.google.com/webstore/detail/wwwyzerdd-for-wikidata/gfidgfgndnaalpihbdjnfkbfniookc>

¹⁵<https://addons.mozilla.org/en-CA/firefox/addon/wwwyzerdd-for-wikidata/>

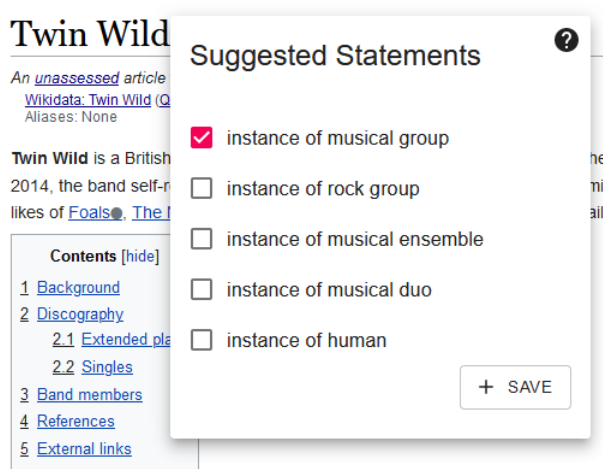


Fig. 5. Psychiq is integrated into Wwyyzzerdd as a drop-down that lets the user select statements to add to Wikidata.




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

A fundamental limitation of Wwyyzzerdd and Psychiq is that they rely on the accuracy and reliability of (English) Wikipedia. The sources that Wwyyzzerdd extracts from Wikipedia are not to the underlying sources used in Wikipedia but to Wikipedia itself. This is better than providing no reference but is not as good as it could get. A future version of the software could allow users to import sources from Wikipedia directly into Wikidata.

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.




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

4.2. Psychiq

In addition to being integrated into Wwyyzzerdd, Psychiq was also converted into a simple “game-like” web interface¹⁶. Figure Table 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. Several days after being announced and ignoring skipped statements, over 97% of suggested statements were accepted. Initial feedback on this tool suggested the model was performing well.

The Psychiq model is robust to missing information. Newly created articles often lack categories and only have titles. For example, for the article “University of Vechta” without categories the model still accurately suggests that it may be a university. The model will not work with only category information because the model assumes the name of the category is the name of the article. For example, just the category “Mosques in Jordan” is predicted to be an instance of “aspect in a geographic region” (Q74817647) when it should have predicted “mosque” (Q32815).

Another use case of Psychiq is identifying errors in Wikidata. By looking for instances where the model is very confident and disagrees with what is currently stated in Wikidata, Psychiq can find probable errors. Unfortunately doing this also produces many uninteresting conflicts where, for instance, the model identifies the item as a road but it’s labeled as a highway. If you filter out uninformative pairs based on their frequency you are left with more actionable ones. Of the remaining top fifty cases where Psychiq is most confident Wikidata is in error¹⁷ manual review found all but one to be actual errors in Wikidata. The most probable detected errors were caused by English Wikipedia pages being converted from disambiguation pages to articles about specific people without updating the corresponding Wikidata item.

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

A major difficulty in Wikidata is establishing and communicating the desired ontology for various classes of information. It is not obvious to new users that “Beat It” (Q210218) should be classed as an instance of “musical work/composition” (Q105543609) and not the more natural sounding “song” (Q7366). Another common confusion is between “bibliography” (Q1631107) and “bibliography” (Q134995) which are two different concepts with the exact same name. Today this information is communicated through Wikiprojects which write up

¹⁶<https://wikidata-game.toolforge.org/distributed/#game=84>

¹⁷<https://www.wikidata.org/wiki/User:BrokenSegue/PsychiqConflicts>

documentation¹⁸ or through model items¹⁹ which are maintained examples of the desired way to model information. Instead, with Psychiq, the correct way to model an item can be suggested using machine learning. If the community wants to change how to represent something, they just needed to update all the existing items to the new ontology and then the next time the model is retrained the suggestions will automatically reflect the new ontological model.

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

5. Conclusion

This paper has discussed two related approaches for improving the user-experience of Wikidata completion using the textual content of Wikipedia. This work represents a step beyond the current dominant tools in the Wikidata 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] 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.
- [9] N. Heist and H. Paulheim, Information extraction from co-occurring similar entities, in: *Proceedings of the Web Conference 2021*, 2021, pp. 3999–4009.
- [10] M. Peel, Pi Bot. <https://github.com/mpeel/wikicode>.
- [11] A. Sarabadani, Kian. <https://github.com/Ladsgroup/Kian>.
- [12] Amir Sarabadani and Daniel Erenrich, Re: Kian, 2023.
- [13] M. Héder and P. Mendes, Round-trip semantics with Sztakipedia and DBpedia Spotlight (2012). doi:10.1145/2187980.2188048.
- [14] 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.
- [15] B. Kratzwald, G. Kunpeng, 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>.
- [16] 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).

¹⁸https://www.wikidata.org/wiki/Wikidata:WikiProject_Music

¹⁹https://www.wikidata.org/wiki/Wikidata:Model_items

- [17] 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>.
- [18] 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.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51