# Human Associations and DBpedia
## Don't Discard Wikilinks

Jörn Hees, Rouven Bauer, Joachim Folz, Damian Borth, and Andreas Dengel

[1] Computer Science Department, University of Kaiserslautern, Germany
[2] Knowledge Management Department, DFKI GmbH, Kaiserslautern, Germany
{firstname.lastname}@dfki.de

**Abstract.** While the amount of Linked Data published in the past years is astonishing and can be seen as a first step towards strong AI, associations as one of the key ingredients of human intelligence and thinking did not receive much attention in the Semantic Web community.
In this paper we describe a semi-automatic mapping approach of strong textual associations from the Edinburgh Associative Thesaurus to DBpedia Entities in the centre of the LOD Cloud. We provide insights into types of associations that can and those that cannot be mapped. After analysing distances and linkage patterns of the resulting associated semantic entities, we conclude that often discarded weak semantic links, such as page-links between Wikipedia articles, contain crucial information. However, utilizing these weak links poses challenges which are detailed as well, along with suggestions on how to work around them.
The generated mapping corresponding to over 25000 human associations is made publicly available and can be used as a benchmark for cross-type link prediction and pattern learning.

## 1 Introduction

In recent years, the Semantic Web [1] and Linked Data [2] have brought us many machine readable datasets and interlinked knowledge bases, and can by this be seen as a first step towards strong AI.

At the same time, associations as one of the building blocks of human intelligence, thinking, context forming and everyday communication [6] are not well represented in the published datasets. This impedes AI research: as a ground truth of semantic entities which are associated by humans is missing, we can neither analyse human associations in existing datasets, nor train machines to learn patterns for them.

Hence, in this paper we present a first human association RDF dataset by transforming the Edinburgh Associative Thesaurus [10] of over 788 K free-text associations into RDF triples [12] in Section 3.

Furthermore, we define the concept of a *semantic association* and provide a first mapping of strong EAT associations to a semantic association of DBpedia [3] entities in Section 4.

This allows us to conduct a first analysis of linkage patterns of human associations in DBpedia/Wikipedia in Section 5 before concluding this paper in Section 6.

## 2 Related Work

To the best of our knowledge the only works towards generating a human association ground truth RDF dataset have been our own. Previously, we developed semantic games with a purpose to collect a semantic association ground truth (Linked Data Games [9], KnowledgeTestGame [7]) or to rank existing triples by association strengths (BetterRelations [8]). While these works can help collecting new associations, the datasets generated in this paper are orders of magnitude larger.

Along the lines of fact ranking ground truth datasets, several other works such as WhoKnows [13] and more recently FRanCo [4] have been published. While fact ranking in general compared to associations focuses only on existing facts, FRanCo in its first step also collected free-text fact input about the entity in question, e.g. "Please tell us the most important facts about Munich". While the question formulation is certainly useful for FRanCo's task to check for missing facts, sadly it is questionable for collecting unbiased associations in the sense of this paper. Nevertheless, it is conceptually the closest of which we are aware to the datasets and mappings described in this paper. In contrast to the mappings presented in this paper, the published NER mapping of the free-text facts back to semantic entities[3] does not seem to have been manually verified and does not seem to be included in the published RDF dataset, yet.

Many other works exist which estimate semantic relatedness, similarity, or associations, for example by textual co-occurrences. We want to point out that this is not in the scope of this paper, as we focus on providing a ground truth dataset that is directly collected from humans. Also we clearly distinguish between associations and other forms of relatedness such as similarity. Two entities can be strongly associated but not similar ("waist - belt") or vice versa ("chair - bed"). In this work we only focus on human associations.

## 3 Edinburgh Associative Thesaurus

After mentioning other related work in the previous section, in this section we will briefly describe the Edinburgh Associative Thesaurus (EAT) dataset and introduce some terminology before detailing our RDF Version of EAT.

EAT [10] was created in the 1970s and is a dataset of single free-text associations collected directly from humans. The associations were collected in several rounds, starting from a seed list of common words as stimuli. In the following rounds, frequent responses of the previous rounds became the new stimuli. Over all rounds, a total of $\sim$ 8200 stimuli were each presented to 100 participants

[3] `http://s16a.org/node/13`

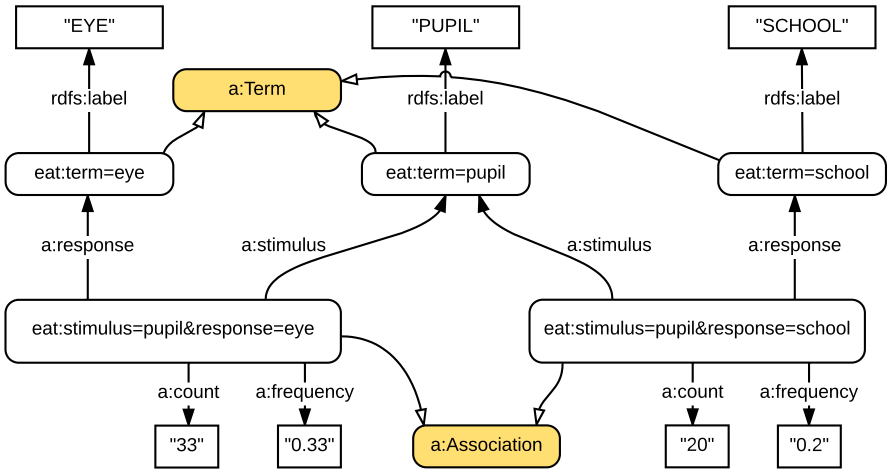


**Fig. 1.** Example of the two EAT associations "pupil - school" (right) and "pupil - eye" (left) as RDF. Notice how the stimulus "pupil" changes its meaning in the two associations from the student (right) to a part of the eye (left).

Prefixes: a: <http://associations.joernhees.de/association_vocab#>, eat: <http://www.eat.rl.ac.uk/#>

(mainly students from Edinburgh) in randomized order. Each participant was allowed to write down one response to each stimulus.

By this, Kiss *et al.* managed to create a well connected network of $\sim 788$ K *raw associations* which form $\sim 326$ K *unique associations* (unique stimulus-response-pairs) between 8200 unique stimuli and $\sim 22700$ unique responses.

About $\sim 5000$ unique associations occur more than 20 times (167 K raw associations). In the remainder of this paper we will refer to them as *strong associations*. An example for such a strong association is the one between stimulus "dog" and response "cat" which occurred 57 out of 100 times.

### 3.1 Association Vocabulary and RDF Version of EAT

The original EAT dataset can be downloaded on the project's website[4], but sadly isn't available as RDF. Hence, we will now describe how the dataset can be transformed into RDF.

We can formally model EAT as a multi-set of raw associations. Each raw association $a \in EAT$ is a free-text stimulus-response-pair: $a = (s, r), s \in S, r \in R$. The union of all stimuli $S$ and responses $R$ forms the set of terms $T = S \cup R$. While the original EAT corpus capitalizes all terms, we will write an association $(\text{PUPIL}, \text{EYE})$ as "pupil - eye" for better readability.

[4] `http://www.eat.rl.ac.uk/`

Further, we can define the count $c_{s,r}$ as the number of occurrences of the raw association

$$c_{s,r} = |\{(s, r) \in EAT\}|$$

and the relative frequency $f_{s,r}$ as the relative count of response $r$ with respect to a fixed stimulus $s$ over all responses to that stimulus:

$$f_{s,r} = c_{s,r} / \sum_{r' \in R} c_{s,r'}.$$

An example for the transformation of the two associations "pupil - eye" and "pupil - school" into RDF can be found in Figure 1.

As EAT consists of free-text associations, we modelled each of its terms $t$ as an RDF literal, keeping its capitalization as found in the original dataset. Further, as RDF does not allow making statements about literals in the subject position, we also mined a URI for each such literal pointing back to the original project's website, for example `eat:term=eye` as shown in Figure 1. This will for example allow us to add additional labels (e.g., other capitalization) to the terms in the future.

Similar to the terms, we also mined a URI pointing back to the original project's website for each unique association $(s, r) \in EAT$, for example `eat:stimulus=pupil&response=eye` and linked its corresponding stimulus, response, count and frequency with the properties defined in our association vocabulary `a`: <`http://associations.joernhees.de/association_vocab#`>.

Further, we assert that each term is a `a:Term` and each association an `a:Association`.

The resulting transformation of EAT into RDF consists of 1 674 376 triples and is provided independently of the following as free download[5].

## 4 Mapping EAT to DBpedia

After the transformation of the EAT dataset into RDF in the previous section, this section describes the process of mapping associations from EAT to equivalent semantic associations between DBpedia entities.

More precisely, we want to find a mapping of each of the terms of an EAT association to two different semantic entities in the DBpedia, such that they convey the same meaning. If we find such two entities, we call the relation between them a *semantic association*.

For example, let's focus on the association "pupil - eye", with URI `eat:stimulus=pupil&response=eye` in Figure 2. We can identify two DBpedia entities, namely `dbp:Pupil` and `dbp:Eye` with the intended meaning of the association and create a new semantic association `dbpam:pupil/eye` with the corresponding links as further detailed in Section 4.3. However, in the association "pupil - school" we find that 'pupil" conveys a different meaning. Instead of a part of the eye it is used synonymous with "student". In this context `dbp:Pupil` conveys a different

[5] `http://associations.joernhees.de/eat.nt.gz`

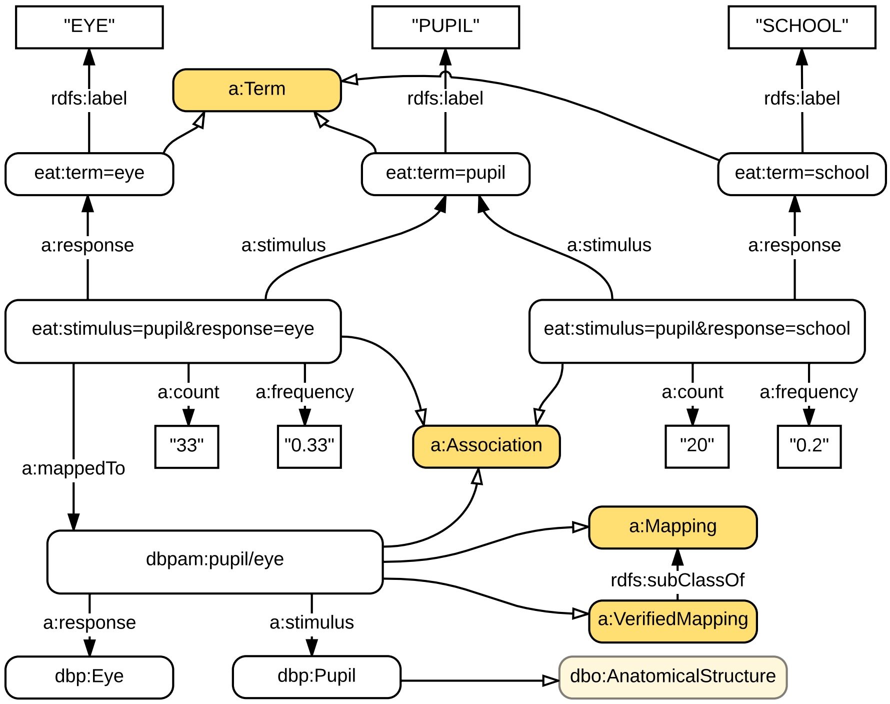


**Fig. 2.** Example mapping from EAT to DBpedia for the association "pupil - eye" (left). As `dbp:Pupil` conveys a different meaning than the association "pupil - school" (right), we refrained from creating a wrong semantic association and mapping.

Prefix: dbp: <http://dbpedia.org/resource/>, dbpam: <http://associations.joernhees.de/mapping_eat_dbpedia#>

meaning than intended in the original association, hence there is no semantic association to `dbp:School` and no mapping. Yet another negative example would be an association "new - york", which forms a composite phrase and does not describe a relation between two semantic entities, and hence is not a semantic association.

For the mapping we will mainly focus on the $\sim 5000$ unique strong associations occurring more than 20 times (167 K raw associations), as they are more robust with respect to subjectivity, location and time dependency.

In the following we will first describe some systematic mapping challenges we identified, setting the expectations for our semi-automatic mapping approach detailed afterwards.

### 4.1 Expected Quantities and Identified Challenges

As the two examples above already show, the mapping process is not straightforward for at least some of the associations. In order to estimate what could be expected from a completely manual mapping, which would involve a lot of human work, we decided to randomly sample 50 out of the $\sim 5000$ unique strong associations and asked two test persons to manually map the stimuli and responses to their corresponding Wikipedia Articles.

The somewhat surprising outcome of this small experiment was that the test persons were only able to manually map 14 of the 50 associations to corresponding semantic associations between DBpedia entities. Out of these 14 the testers reported that 6 could be matched following very simple rules. The remaining 8 required human knowledge and understanding, for example to pick a synonymous term from a list of alternatives on a disambiguation page.

Because of this, our expectation for any at least partly automated mapping process is that we will only be able to achieve a successful mapping in about 6/50 to 14/50 of the cases, as even humans cannot do better. Starting from $\sim 167$ K raw associations, this means that we can expect to generate a mapping for 20 K to 46.8 K of them.

We also asked the testers to collect notes about the associations they could not map or had difficulties with. The notes can be summarized as follows (including overlaps):

- **Composite phrases**: In 12/50 cases the association formed a composite phrase (e.g., "identical - twins"), which is just a single semantic entity in DBpedia (not a semantic association).
- **Synonyms**: In 9/50 cases the stimulus and response were synonyms (e.g., "children - kids") leading to the same semantic entity in DBpedia (not a semantic association).
- **Adjectives / Verbs**: In 11/50 cases at least one of the terms was an adjective (e.g., "hot - cold"), in 3/50 one was a verb (e.g., "ring - bell"). As Wikipedia is an encyclopaedia, there is a bias towards substantives, often making it harder to indisputably map adjectives or verbs to a semantic entity in DBpedia.
- **False friends**: In 4/50 cases a simple lookup of the stimulus or response works, but leads to a wrong entity (e.g., "sharpen - knife"[6]).

Further inspection revealed that in 8/50 cases one of the terms was a plural word (e.g., "colours - red"), which can mostly be handled without problems due to the existence of Wikipedia redirect pages.

### 4.2 Semi-Automatic Mapping Approach

Using the observations from the previous section we developed the semi-automatic mapping approach described in this section. Our approach aims to find high

[6] "Sharpen" describes an Eclipse (IDE) plugin: `http://en.wikipedia.org/w/index.php?title=Sharpen&oldid=629433221`

quality mappings from strong EAT association to semantic associations between DBpedia entities, while reducing the amount of necessary human work.

In order to achieve this, we use a two step process: First, we perform an automatic mapping, employing a scoring component which focuses on the identified mapping challenges. Afterwards, we let humans verify the highest scoring mappings with a web application to guarantee high precision of the generated mappings.

**Automatic Mapping with Scoring Component** The automatic mapping uses the Wikipedia API[7] to perform simple searches (following redirect pages) for the stimulus and response in article titles and full texts in order to generate candidate mappings.

The scoring component then rates these candidate mappings, mostly by trying to identify the potential problems mentioned in Section 4.1, helping us to focus on the least disputable mapping candidates first:

- **Composite phrases** (e.g., "port - wine"): As a composite phrase is a name for a single semantic entity it is a bad candidate for a semantic association (between two different semantic entities). Hence, if searching for Wikipedia articles (or redirect pages) containing stimulus and response in their title is successful, the mapping's score receives a strong punishment.
- **Reflexive mappings / synonyms** (e.g., "child - children"): If the mapping of both the stimulus and the response result in the same semantic entity, the score is strongly punished.
- **Adjectives & verbs vs. nouns** (e.g., "unbound - free"): Due to Wikipedia's nature of being an encyclopaedia, adjectives and verbs are under-represented in contrast to nouns. To identify such cases, the stimulus and response are searched in Wordnet [5], potentially resulting in multiple synset candidates for each. Mappings containing only synset candidates with the given type "noun" are slightly rewarded. The more synset candidates with types unequal to "noun" are found, the stronger the punishment for the mapping's score.
- **Plural words** (e.g., "thumbs - fingers"): A simple stemming approach is used to compare the stimulus/response to the identified Wikipedia article titles after following redirects. If the match is close to perfect and only differs in singular/plural, the score only receives a slight punishment.
- **Disambiguation pages** (e.g., "pod - pea"): If the mappings of stimulus or response result in a Wikipedia disambiguation page, the mapping's score receives a strong punishment.

After applying the automatic mapping component with this scoring mechanism to the $\sim$ 5000 strong associations, 1066 semantic association candidates (corresponding to $\sim$ 34.2 K raw associations) remained for human verification.

[7] `http://www.mediawiki.org/wiki/API:Main_page`

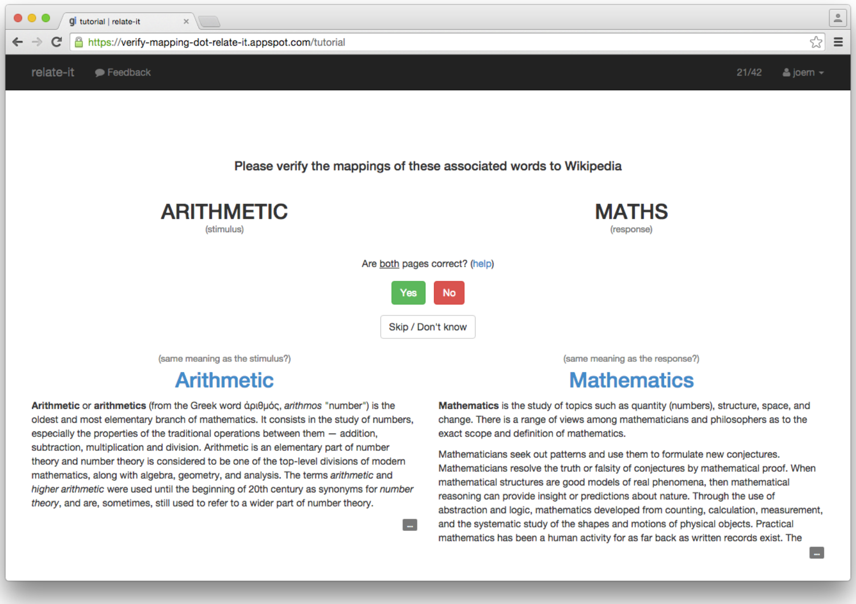


**Fig. 3.** DBpedia Mapping Verification Web Application. On top the original EAT association can be seen. Below the two identified corresponding Wikipedia articles and their abstracts are shown.

**Manual Verification** In order to quickly verify the mapping candidates from the previous section, we developed a small mapping verification web application that shows the textual association from EAT on top (stimulus - response) and both mapped Wikipedia articles below (featuring their abstracts). After a tutorial explaining the purpose, the user is asked if both stimulus and response are correctly mapped to Wikipedia pages. Possible answers are "Yes", "No" or "Skip / Don't know", as can be seen in Figure 3.

The mapping candidates are presented in randomized order and at most once to each user. Candidates that receive a "No" or "Skip / Don't know" rating are immediately excluded from further verifications. After receiving three "Yes" ratings from different users, a mapping candidate is marked as "valid" and excluded from further evaluation in order to concentrate verification on the remaining candidates.

The web application was used by 10 reviewers and quickly allowed the verification of 790 of 1066 mappings (corresponding to $\sim 25.5$ K raw associations), which is well within the expected quantities mentioned in Section 4.1.

### 4.3 Mapping Results as RDF Dataset

For each of the 790 verified mapped associations an additional mapping URI is created in the `dbpam`: <`http://associations.joernhees.de/mapping_eat_` dbpedia#> name space (e.g., dbpam:pupil/eye in Figure 2) and linked from the EAT association with the a:mappedTo property. The mapping URI is also linked to the DBpedia stimulus and response accordingly, as well as typed as an a:Association, a a:Mapping and a a:VerifiedMapping.

The resulting mapping dataset consisting of 4740 triples can be downloaded[8] or simply dereferenced.

## 5 Analysis of Semantic Associations in DBpedia

Based on the 790 manually verified mappings, we first provide some statistics about the mapping and involved nodes, before performing a first analysis of the distances and linkage patterns in DBpedia.

727 out of the 790 mappings lead to a distinct DBpedia stimulus-response-pair. For example, the three associations "casks - beer", "barrels - beer" and "barrel - beer" all lead to the DBpedia stimulus dbp:Barrel and response dbp:Beer. In order not to skew our results towards such multiple mappings, we will focus on the distinct 727 stimulus-response-pairs of DBpedia entities in the following and call them DBpedia associations for brevity.

Out of these 727 DBpedia associations, there are 685 distinct stimulus and 346 distinct response nodes, totalling in 955 distinct nodes. None of the stimuli occur more than twice, but some of the responses occur more frequently, such as dbp:Money or dbp:Bird, as can be seen in Table 1.

**Table 1.** Most frequent response nodes

| Response | Count | Response | Count |
|---|---|---|---|
| dbp:Money | 19 | dbp:Water | 9 |
| dbp:Bird | 15 | dbp:Army | 8 |
| dbp:Horse | 14 | dbp:Beer | 8 |
| dbp:Automobile | 13 | dbp:Death | 7 |
| dbp:Flower | 12 | dbp:Fish | 7 |
| dbp:Music | 12 | dbp:Bed | 7 |
| dbp:Tree | 11 | dbp:Ship | 7 |
| dbp:Sea | 11 | dbp:Red | 6 |
| dbp:Dog | 9 | dbp:Gun | 6 |
| dbp:Food | 9 | dbp:Hair | 6 |

For our analysis of the distances and linkage patterns in DBpedia, we used a local Virtuoso 7.2[9] mirror of the DBpedia 2015-04[10] core and extended datasets. The *core* dataset[11] includes the $\sim 412$ M triples which are loaded on the public

[8] http://associations/joernhees.de/mapping_eat_dbpedia.nt.gz
[9] https://github.com/openlink/virtuoso-opensource
[10] http://wiki.dbpedia.org/dbpedia-data-set-2015-04
[11] http://downloads.dbpedia.org/2015-04/core/

DBpedia SPARQL endpoint[12]. Additionally, we extended our endpoint with all datasets for the English DBpedia[13] which were not already included in core and the nearly 159 M Wikilinks (`dbo:wikiPageWikiLink`), to which we refer as the *extended* dataset.

In order to analyse the differences between the core and extended dataset, we first computed the degrees of all DBpedia association nodes. As expected, the node degrees in the extended dataset are much larger than the ones in the core dataset (avg. $\sim$ 4650 extended vs. $\sim$ 1240), as can be seen in Table 2. Nevertheless, we can observe, that even without Wikilinks, some of the nodes, such as `dbp:Animal`, `dbp:Insect`, `dbp:France` have a very high degree. Investigations revealed that such high node degrees are mostly originating from incoming edges such as `dbo:kingdom`, `dbo:class`, `dbo:country`, `dbo:type`, `dbo:order` and `dbo:birthPlace` in the core dataset. In the extended dataset, they unsurprisingly mostly originate from incoming edges of the property `dbo:wikiPageWikiLink`, but also from `gold:hypernym` from the Linked Hypernym Datasets [11].

**Table 2.** Top-20 degrees of the 955 investigated association nodes in the core (left) and extended (right) datasets.

| Node | Degree | Node | Degree |
|---|---|---|---|
| dbp:Animal | 237855 | dbp:Animal | 445324 |
| dbp:Insect | 118589 | dbp:Village | 344264 |
| dbp:France | 94826 | dbp:Insect | 239032 |
| dbp:India | 85386 | dbp:France | 234700 |
| dbp:Plant | 79062 | dbp:India | 196686 |
| dbp:Italy | 55966 | dbp:Plant | 149369 |
| dbp:Village | 54082 | dbp:Italy | 143942 |
| dbp:Beetle | 43739 | dbp:Town | 85994 |
| dbp:Scotland | 27607 | dbp:Beetle | 83109 |
| dbp:Bird | 25933 | dbp:Scotland | 73312 |
| dbp:Switzerland | 19874 | dbp:Paris | 66504 |
| dbp:City | 18030 | dbp:Switzerland | 61214 |
| dbp:Paris | 17362 | dbp:City | 53008 |
| dbp:Wales | 14605 | dbp:Bird | 50332 |
| dbp:Town | 13301 | dbp:Ireland | 40592 |
| dbp:Ireland | 11340 | dbp:Marriage | 38643 |
| dbp:Rome | 10344 | dbp:Rome | 38611 |
| dbp:Fly | 10299 | dbp:Wales | 38532 |
| dbp:Mayor | 9812 | dbp:School | 32824 |
| dbp:Reptile | 9595 | dbp:Novel | 32193 |

Next, we analysed the minimum path lengths between stimulus and response of the DBpedia associations. In the core dataset, only 34 ($< 5\%$) of the 727

[12] `http://dbpedia.org/sparql`

[13] `http://downloads.dbpedia.org/2015-04/core-i18n/en/`

DBpedia associations are directly connected (24 forward, 12 backward, 2 bi-directionally) and still only 417 (57.4%) via another node (path of length 2). In contrast to this, in the extended dataset 547 (75.2%) of the 727 DBpedia associations are directly connected (445 forward, 413 backward and even 311 bi-directionally) and 726 (99.9%) via another node (path of length 2).

For paths of length 1, we also analysed which properties frequently link the stimulus and response nodes. In the core dataset for the 34 associations these properties are mostly `rdfs:seeAlso`, `dbo:class`, `dbprop:classis`, `dbo:kingdom`, `dbo:country`, and `dbo:ingredient` (unidirectional). In the extended dataset for the 547 associations, we additionally find many `dbo:wikiPageWikiLink` and `gold:hypernym`.

Furthermore, in more than 60% an existing Wikilink connecting a DBpedia association is bi-directional. As globally only $\sim$ 7% of all Wikilinks are bi-directional, we seem to have identified a strong signal for a semantic association from the dataset of Wikilinks, which is otherwise difficult to use due to its quantity and weak semantics.

Finally, we also analysed the properties and connecting nodes for paths of length 2. In the core dataset, the majority of connecting properties consists of `dcterms:subject`, `rdf:type`, `rdfs:seeAlso`, `dbo:product` and `dbo:class`. Connecting nodes are unsurprisingly `owl:Thing`, but also nodes such as `umbel:EukaryoticCell`, `umbel:BiologicalLivingObject`, `umbel:Animal`, `umbel:Bird` and `dbp:Category:Plant_morphology`. In the extended dataset the connecting properties are again additionally lead by `dbo:wikiPageWikiLink` and `gold:hypernym`, followed by `dbprop:wikiPageUsesTemplate`. The connecting nodes are additionally lead by nodes such as `dbp:Template:Reflist`, `dbo:Article`, but also `dbp:QI_(L_series)`, `dbp:List_of_Latin_words_with_English_derivatives` and `dbp:Bird`.

Again, we can see that a lot of information is hidden within the Wikilinks. For example, connecting nodes such as `dbp:QI_(L_series)`[14] link to very many common words and manage to connect many of the DBpedia association nodes in the extended dataset, but fail to do so in the core dataset without Wikilinks.

## 6 Conclusion & Outlook

In this paper we presented a transformation of 788 K free-text associations from the Edinburgh Associative Thesaurus into a RDF dataset, making it easily accessible and mappable to other datasets in the Linked Data community.

Further, we provide a first such mapping to semantic associations between DBpedia entities. We described our semi-automatic approach to find indisputable, high precision mappings, resulting in 790 manually verified mappings corresponding to $\sim$ 25.5 K raw associations, and which challenges were identified in the process.

With the generated datasets we conducted a first analysis of the distances and linkage patterns of semantic associations in the DBpedia. We note signifi-

[14] QI is a game show featuring many common words.

cant differences between the DBpedia Core and Extended datasets, mainly with respect to the Wikipedia page links.

In the future we plan to conduct further analysis and pattern learning based on the mapped semantic associations. As all generated datasets are publicly available, we also look forward to them being used as benchmark or ground truth datasets, for example for link prediction tasks.

This work was financed by the University of Kaiserslautern PhD scholarship program and the BMBF project MOM (Grant 01IW15002).

## References

1. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. Scientific American 284(5), 34–43 (may 2001)
2. Bizer, C., Heath, T., Berners-Lee, T.: Linked Data - The Story So Far. International Journal on Semantic Web and Information Systems 5(3), 1–22 (jan 2009)
3. Bizer, C., Lehmann, J., Kobilarov, G., Auer, S., Becker, C., Cyganiak, R., Hellmann, S.: DBpedia - A crystallization point for the Web of Data. Web Semantics: Science, Services and Agents on the World Wide Web 7(3), 154–165 (sep 2009)
4. Bobić, T., Waitelonis, J., Sack, H.: FRanCo – A Ground Truth Corpus for Fact Ranking Evaluation. In: SumPre 2015 - 1st International Workshop on Summarizing and Presenting Entities and Ontologies @ ESWC 2015. pp. 1–12 (2015)
5. Fellbaum, C. (ed.): WordNet: An Electronic Lexical Database. MIT Press, Cambridge, MA, USA (1998)
6. Gerrig, R.J., Zimbardo, P.G.: Psychology and Life. Allyn & Bacon, Pearson, Boston, USA, 19th edn. (2010)
7. Hees, J., Khamis, M., Biedert, R., Abdennadher, S., Dengel, A.: Collecting Links between Entities Ranked by Human Association Strengths. In: ESWC. vol. 7882, pp. 517–531. Springer LNCS, Montpellier, France (2013)
8. Hees, J., Roth-berghofer, T., Biedert, R., Adrian, B., Dengel, A.: BetterRelations: Using a Game to Rate Linked Data Triples. In: KI 2011: Advances in Artificial Intelligence. pp. 134–138. Springer Berlin / Heidelberg, Berlin (2011)
9. Hees, J., Roth-Berghofer, T., Dengel, A.: Linked Data Games: Simulating Human Association with Linked Data. In: LWA 2010. pp. 255–260. Kassel, Germany (2010)
10. Kiss, G.R., Armstrong, C., Milroy, R., Piper, J.: An associative thesaurus of English and its computer analysis. In: Aitken, A.J., Bailey, R.W., Hamilton-Smith, N. (eds.) The Computer and Literary Studies, pp. 153–165. Edinburgh University Press, Edinburgh, UK (1973)
11. Kliegr, T.: Linked hypernyms: Enriching DBpedia with Targeted Hypernym Discovery. Web Semantics: Science, Services and Agents on the World Wide Web 31, 59–69 (2015)
12. Klyne, G., Carroll, J.J.: Resource Description Framework (RDF): Concepts and Abstract Syntax. Tech. rep., World Wide Web Consortium (feb 2004)
13. Kny, E., Kölle, S., Töpper, G., Wittmers, E.: WhoKnows? (oct 2010)