

CounQER: A System for Discovering and Linking Count Information in Knowledge Bases

Shrestha Ghosh¹, Simon Razniewski¹, and Gerhard Weikum¹

Max Planck Institute for Informatics, Saarland Informatics Campus, 66123
Saarbruecken, Germany
{ghoshs, srazniew, weikum}@mpi-inf.mpg.de

Abstract. Predicate constraints of general-purpose knowledge bases (KBs) like Wikidata, DBpedia and Freebase are often limited to subproperty, domain and range constraints. In this demo we showcase CounQER, a system that illustrates the alignment of *counting predicates*, like `staffSize`, and *enumerating predicates*, like `workInstitution`⁻¹. In the demonstration session, attendees can inspect these alignments, and will learn about the importance of these alignments for KB question answering and curation. CounQER is available at <https://counqer.mpi-inf.mpg.de/spo>.

Keywords: Knowledge bases · semantics · count information.

1 Introduction

Motivation and problem. Detecting inter-predicate relations in Knowledge Bases (KBs) beyond inheritance can lead to a better semantic understanding that can be leveraged for important tasks such as KB curation and question answering (QA). In this work we focus on set predicates and their alignment. Set predicates describe the relation between an entity and a set of entities through two variants - i) *counting predicates* which relate an entity to a count (of a set of other entities) and, ii) *enumerating predicates* which relate an entity to multiple entities.

Consider a list of counting predicates, `{numberOfChildren, staffSize}`, which take only integer count values as objects and a list of enumerating predicates, `{child, employer`⁻¹, `workInstitution`⁻¹}, which take only entity values. Identifying set predicates pairs across the two variants, aligned by their semantic relatedness, such as `{numberOfChildren ↔ child}`, `{staffSize ↔ employer`⁻¹}, `{staffSize ↔ workInstitution`⁻¹}, has two major benefits.

1. *KB curation* - We can discover incompleteness and/or inconsistencies in KBs through alignments [7,9]. For instance, if the value of `numberOfChildren` exceeds the count of `child` entities of a subject, then the `child` statements for that subject may be incomplete. Alternately, if the value of `numberOfChildren` is less than the count of `child` entities, there may be inconsistent enumerations. An empty instantiation is also an indication of incompleteness.

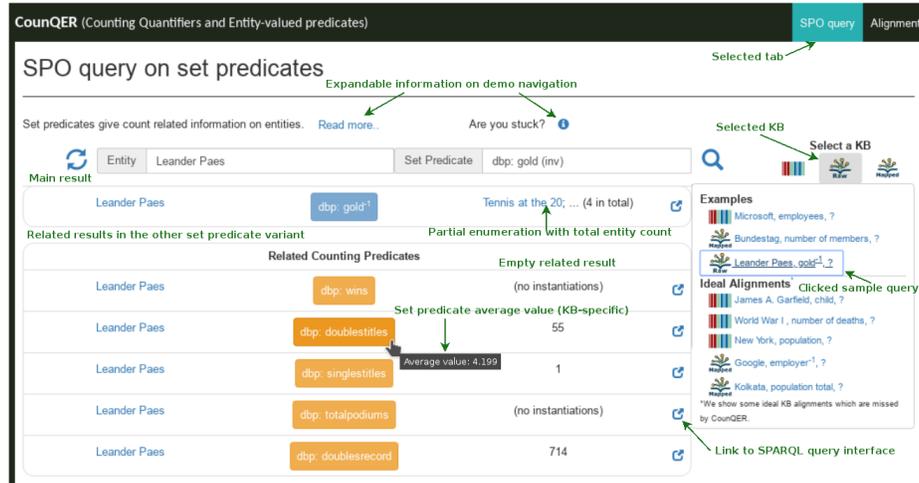


Fig. 1: The interface for SPO queries showing results on an example query.

2. *QA enhancement* - Set predicate alignments can aid in KB query result debugging and enrichment [1,2]. Even in an event of empty result, for instance, when an entity has no `numberOfChildren` predicate instantiations, but, has `child` predicate instances, we can enumerate the object entities of `child` instead.

Set predicate alignments highlight the variation in predicate usage for the same concept. For instance, the `staffSize` of an entity has related results on employees through `employer-1` as well as `workInstitution-1`.

Approach. CounQER (short for “**C**ounting **Q**uantifiers and **E**ntity-valued **P**redicates”) uses a two-step approach. First it identifies the counting and enumerating predicates with supervised classification and then aligns set predicate pairs, one from each variant, according to ranking methods and statistical and lexical metrics. For further details refer to [3]. The classification and alignment steps are executed offline. We use the obtained results in our demonstrator for count-related SPO queries on three KBs¹ - Wikidata-truthy and two variants of DBpedia based on mapped and raw extractions.

Fig. 1 shows the interface with results on an example query on the DBpedia-raw KB. The query is on the events where the entity, *Leander Paes*, wins gold (`dbp: gold-1`). The main result (set predicate in blue) is succeeded by related results on ranked and aligned set predicates (in orange). Enumerations expand on hovering and we show up to 1000 enumerations. A user can check the actual query fired for each row by following the link to the SPARQL query endpoint. Also, KB-specific predicate statistics show up while hovering over the predicate

¹ <https://tinyurl.com/wikidata-truthy>, <https://tinyurl.com/dbpedia-mappings>, <https://tinyurl.com/dbpedia-raw>

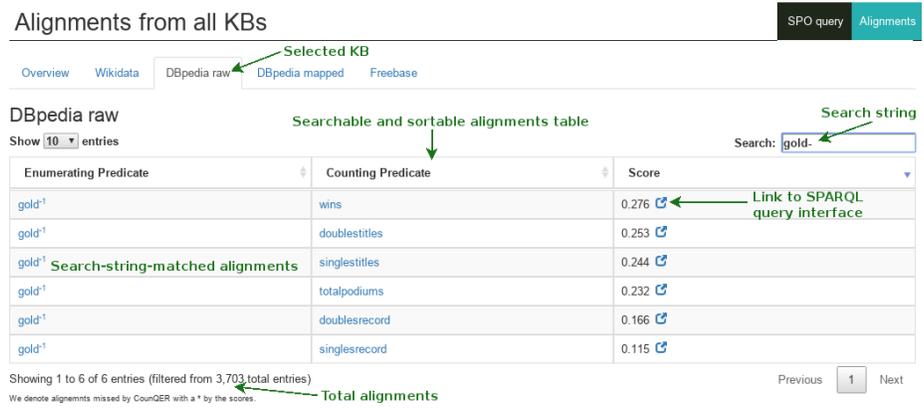


Fig. 2: Interface for viewing KB-specific alignments.

buttons. On clicking a set predicate from the related results, a new query is fired on the same subject and the new clicked predicate.

The complete ranked list of set predicate alignments for the three KBs as well as Freebase can be viewed as in Fig. 2. Here too, we provide links to the SPARQL endpoint showing the subjects that have populated facts for the alignments.

Related Work. Schema and ontology alignment is a classic problem in data integration, which in the semantic web is tackled by focusing on the dynamics of entity relations across ontologies [8]. Subset and equivalence relation alignment is one of the popular approaches to ontology alignment [4]. The problem of aligning enumerations with counts, which we address is atypical since most approaches do not target completeness and correctness of KBs [6]. Even though between 5% to 10% of questions in popular TREC QA datasets deal with counts [5], QA systems like AQQU [1] and QAnswer [2] only perform ad-hoc count aggregation function to deal with typical count questions, which start with “How many..?”.

2 System Description

SPO query. The SPO query function provides two input fields, *Entity* and *Set Predicate*, and a KB selection button. The first field provides real-time entity suggestions from the selected KB, based on the input prefix, to the user to choose from. Next, the user selects a set predicate from the set predicate input field. The predicate choices are KB-specific and ordered by i) whether they are populated and have alignments, ii) they are populated but without alignments, and iii) they are unpopulated.

Upon execution, the input parameters are sent to our server, where we determine the variant of the user-selected set predicate - counting or enumerating. Then from the KB-specific alignments containing the queried predicate, we shortlist the top-five highest scoring pairs to obtain related set predicate facts.

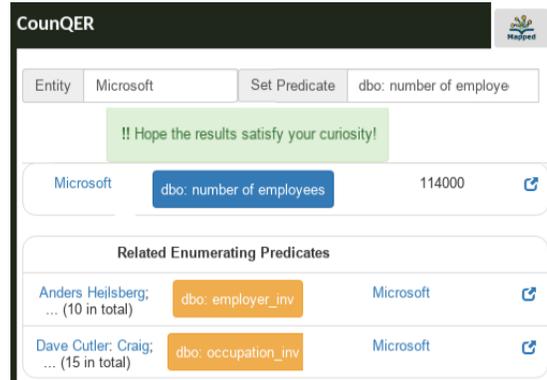


Fig.3: Query results on the number of employees in Microsoft (DBpedia-mapped).

If there are no alignments we do not generate any related query. The server then fires the main query to the SPARQL endpoint of the corresponding KB followed by the SPARQL queries for the aligned set predicates, if present. Once these results are obtained, the server returns the results along with KB-specific predicate statistics, *i.e.*, the average value that the counting predicates take and the average number of entities per subject that the enumerating predicates take.

Alignments. CounQER provides an option of viewing all alignments across the four KBs along with their alignment scores. A user can go through the list ordered by the alignment score or, use the search bar to filter matching set predicates and view their corresponding alignments. Each alignment has a link to SPARQL query API where the user can view the list of subjects for which the predicate pair co-occur.

The main features of the interface are as follows.

1. *Predicate suggestions* - Set predicates are ordered based on whether they are populated for the selected entity and whether alignments exist for them.
2. *Empty results* - If the main query returns an empty result, but, the predicate has populated alignments, CounQER shows the related results. Conversely, if the set predicate in the main query is populated and alignments exist for this predicate, we show the related results regardless of them being empty, thus highlighting potential incompleteness in the KB, w.r.t the queried entity.
3. *Links to SPARQL queries* - Every row in the results contains a link to the SPARQL endpoint, which a user can follow to check the actual query that was fired and also view enumerations of size more than 1000. Alignment tables also link to the SPARQL endpoint with queries which list subjects for which the set predicate pair co-occur.

We also show some manually added ideal alignments, *i.e.*, the alignments which are present in the investigated KBs but missed by the automated CounQER methodology. These alignments are also present in the table with a fictitious score between [0.9-1].

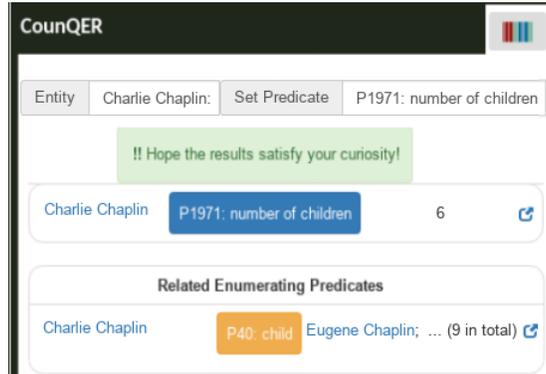


Fig. 4: Query results on the number of children of Charlie Chaplin (Wikidata).

3 Demonstration Experience

Scenario 1 - QA. In a query about the number of employees at Microsoft, CounQER finds the main result from the queried KB, DBpedia-mapped, to be 114,000 employees. In addition, CounQER returns instantiated facts on interesting enumerating predicates, such as, employer^{-1} and occupation^{-1} (see Fig. 3).

Scenario 2 - KB curation. Consider the example in Fig. 4, where the user searches for the number of children of the British comic actor, Charlie Chaplin. The alignment results reveal inconsistent information in Wikidata-truthy. While the value for `number of children` is 6, there are 9 statements for the enumerating predicate `child`.

Next, we investigate the winning titles of Roger Federer in DBpedia-raw (Fig. 5). Even though a query on the golds won by Federer returns no main results, unlike the query on the golds won by Leader Paes in Fig. 1, the counting predicates `doublestitles` (2^{nd}) and `singlestitles` (3^{rd}) give the number of doubles and singles titles won by Federer.

4 Conclusion

We demonstrate how set predicate alignments highlight redundancies in the KB schema, enhance question answering by providing supporting counts and/or enumerations and help in KB curation. Utilizing KB alignments to automatically flag inconsistent SPO facts for resolution and highlight SPO facts needing completions is a possible future work. Analysing multi-hop alignments and extending KB alignments towards open information extraction is also worth exploring.

The screenshot shows the CounQER interface. At the top, there are two input fields: 'Entity' with the value 'Roger Federer' and 'Set Predicate' with the value 'dbp: gold'. Below these fields is a yellow warning box that says '!! There exist no instantiated facts on the queried set predicate.' Underneath is a section titled 'Related Counting Predicates' which contains a table of related predicates for 'Roger Federer'.

Entity	Predicate	Count	Link
Roger Federer	dbp: wins	(no instantiations)	↗
Roger Federer	dbp: doublestitles	8	↗
Roger Federer	dbp: singlestitles	88	↗
Roger Federer	dbp: totalpodiums	(no instantiations)	↗
Roger Federer	dbp: doublesrecord	(no instantiations)	↗

Fig. 5: Query results on golds won by Roger Federer (DBpedia-raw).

References

1. Bast, H., Haussmann, E.: More accurate question answering on Freebase. In: Proceedings of the 24th ACM International Conference on Information and Knowledge Management, CIKM (2015)
2. Diefenbach, D., Migliatti, P.H., Qawasmeh, O., Lully, V., Singh, K., Maret, P.: QAnswer: A question answering prototype bridging the gap between a considerable part of the LOD cloud and end-users. In: The World Wide Web Conference, WWW (2019)
3. Ghosh, S., Razniewski, S., Weikum, G.: Uncovering hidden semantics of set information in knowledge bases. arXiv:2003.03155 (Under revision at Journal of Web Semantics 2020)
4. Koutraki, M., Preda, N., Vodislav, D.: Online relation alignment for linked datasets. In: The Semantic Web - 14th International Conference, ESWC (2017)
5. Mirza, P., Razniewski, S., Darari, F., Weikum, G.: Enriching knowledge bases with counting quantifiers. In: The Semantic Web - 17th International Semantic Web Conference, ISWC (2018)
6. Paulheim, H.: Knowledge graph refinement: A survey of approaches and evaluation methods. Semantic Web (2017)
7. Razniewski, S., Suchanek, F., Nutt, W.: But what do we actually know? In: Proceedings of the 5th Workshop on Automated Knowledge Base Construction AKBC (2016)
8. Shvaiko, P., Euzenat, J.: Ontology matching: state of the art and future. IEEE Transactions on Knowledge and Data Engineering (2013)
9. Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J., Auer, S.: Quality assessment for linked data: A survey. Semantic Web (2016)