

Cardinal Virtues: Extracting Relation Cardinalities from Text

Paramita Mirza¹, Simon Razniewski², Fariz Darari², Gerhard Weikum¹

¹ Max Planck Institute for Informatics

² Free University of Bozen-Bolzano

{paramita, weikum}@mpi-inf.mpg.de

{razniewski, darari}@inf.unibz.it

Abstract

Information extraction (IE) from text has largely focused on relations between individual entities, such as who has won which award. However, some facts are never fully mentioned, and no IE method has perfect recall. Thus, it is beneficial to also tap contents about the cardinalities of these relations, for example, how many awards someone has won. We introduce this novel problem of extracting cardinalities and discuss specific challenges that set it apart from standard IE. We present a distant supervision method using conditional random fields. A preliminary evaluation results in precision between 3% and 55%, depending on the difficulty of relations.

1 Introduction

Motivation Information extraction (IE) can infer relations between named entities from text (e.g., (Mitchell et al., 2015; Del Corro and Gemulla, 2013; Mausam et al., 2012)), yielding for example which awards an athlete has won, or instances of family relations like spouses, children, etc. These methods can be harnessed for summarization, question answering (QA), and more. For populating knowledge bases (KBs), the IE output is usually cast into subject-predicate-object (SPO) triples, such as $\langle BarackObama, hasChild, Malia \rangle$, or sometimes n -ary tuples such as $\langle MichaelPhelps, hasWon, OlympicGold, 200mButterfly, 2016 \rangle$.

IE has focused on capturing full SPO triples (or n -ary facts) with all arguments bound to entities for relation P . However, news, biographies or discussion forums often contain numeric expressions that reveal cardinalities of relations. Phrases such as “her two children” or “his 28th medal” are valu-

able cues for quantifying the *hasChild* and *hasWon* relations. This can be harnessed in QA for cases like “Who won the most Olympic medals?”

An important application of relation cardinalities is KB curation. KBs are notoriously incomplete, contain erroneous triples, and are limited in keeping up with the pace of real-world changes. For example, a KB may contain only 10 of the 28 Olympic medals that Phelps has won, or may incorrectly list 3 children for Obama. Extracting the cardinalities of relations for given subject entities can address all of these issues.

Relation cardinalities are disregarded by virtually all IE methods. Open IE methods (Mausam et al., 2012; Del Corro and Gemulla, 2013) capture triples (or quadruples) such as $\langle Obama, has, two\ children \rangle$. However, there is no way to interpret the numeric expression in the O slot of this triple. While IE methods that hinge on pre-specified relations for KB population (e.g., NELL (Mitchell et al., 2015)) can already capture numeric values for explicitly stated attributes such as $\langle Berlin2016attack, hasNumOfVictims, 32 \rangle$, they are currently not able to learn them.

This paper addresses the novel task of extracting relation cardinalities. For a given subject entity s and predicate p , we aim to infer the cardinality $|\{ \langle S, P, O \rangle \mid S = s, P = p \}|$ directly from text, without having to observe any O entities. This task poses several challenges:

- *IE Training*. Most IE methods build on seed-based distant supervision. However, if the underlying KB is not complete, taking the counts of SPO triples for a given SP pair may result in wrong seeds, which can lead to poor patterns.
- *Compositionality*. The cardinality of an SP pair for a relation may depend on several cardinality mentions. For example, when observing “Angelina has *two* sons and *three* daughters”, one could infer the children cardinality by summing.

- *Linguistic Variance.* In addition to cardinal numbers, cardinality IE should also pay attention to number-related terms, e.g., “Angelina gives birth to *twins*”, or ordinal information, e.g., “Angelina’s *fourth* child”, which can reveal lower bounds on relation cardinalities.

Approach and Contribution Our method learns patterns of phrases that contain cardinal numbers, relying on the distant supervision approach by counting facts for given SP pairs. Our technical contributions are as follows: (i) we provide a statistical analysis of numeric information in Wikipedia articles; (ii) we develop a CRF-based extraction method for relation cardinalities that achieves precision scores of up to 55%; (iii) we analyze further challenges in this research and outline possible solutions.

2 Related Work

Knowledge Bases and Information Extraction

Automated KB construction is a major effort for quite a while. Some approaches, such as YAGO (Suchanek et al., 2007) or DBpedia (Auer et al., 2007), focus on structured parts of Wikipedia, while other approaches such as OLLIE (Mausam et al., 2012), ClauseIE (Del Corro and Gemulla, 2013) or NELL (Mitchell et al., 2015), focus on unstructured contents across the whole Web. In the latter, usually the schema is also not predefined, thus such approaches are called Open IE. Most state-of-the-art systems now rely on distant supervision (Craven and Kumlien, 1999; Mintz et al., 2009).

Despite all efforts, KBs are immensely incomplete. For instance, the average number of children per person in Wikidata (Vrandečić and Krötzsch, 2014) is just 0.02 (Razniewski et al., 2016).

Numbers and Relation Cardinalities Numbers in text are an important source of information. Much work has been done on understanding numbers that express temporal information (Ling and Weld, 2010; Strötgen and Gertz, 2010), and more recently, on numbers that express physical quantities or measures, either mentioned in text (Chaganty and Liang, 2016) or in the context of web tables (Ibrahim et al., 2016; Neumaier et al., 2016).

In contrast, numbers that express relation cardinalities have received little attention so far. State-of-the-art Open-IE systems either hardly extract cardinality information or fail to extract cardinal-

NE Tag	Frequency
DATE, TIME, DURATION, SET (temporal)	54.28%
NUMBER	40.13%
<i>Relation cardinality</i>	18.86%
PERCENT	2.92%
MONEY	2.25%
PERSON, LOCATION, ORGANIZATION	0.26%
ORDINAL	0.16%

Table 1: NE-tags of numbers in Wikipedia.

ities at all. While NELL, for instance, knows 13 relations about the number of casualties and injuries in disasters, they all contain only seed facts and no learned facts. The only prior work we are aware of is of Mirza et al. (2016), who use manually created patterns to mine children cardinalities from Wikipedia. It is shown that with 30 manually crafted patterns and simple filters it is possible to extract 86,227 children-cardinality-assertions with a precision of 94.3%.

3 Relation Cardinalities

Definition We define a mention that expresses relation cardinalities as the following: “A *cardinal number that states the number of objects that stand in a specific relation with a certain subject.*”

Using this definition, we analyzed how often relation cardinalities occur in Wikipedia. Relying on the part-of-speech (PoS) tagger of Stanford CoreNLP (Manning et al., 2014), we extracted numbers—i.e., words tagged as CD (cardinal number)—from 10,000 random Wikipedia articles. The distribution of their named-entity (NE) tags, according to Stanford NE-tagger, is shown in Table 1. While temporal-related numbers are the most frequent, around 40% are classified only as unspecific NUMBER. By manually checking 100 random NUMBERS, we observed that 47 are relation cardinalities,¹ i.e., approximately 18.86% of all numbers in Wikipedia are relation cardinalities.

We also analyzed the nouns frequently modified by NUMBERS, based on their dependency paths, finding *people, games, children, times, members* and *seasons* among the top nouns. Coarse topic-grouping of the nouns shows that most NUMBERS are about sport (*games, goals*), followed by artwork (*seasons, books*), politics and organization (*members, countries*), and family (*children*).

¹ Among the others are measures, age, or expressions like “one of the...”.

4 Relation Cardinality Extraction

Ideally, we would like to make sense of all cardinality statements found in text. However, this would require us to resolve the meaning of a large set of vague predicates, which is in general a difficult task. We thus turn the problem around: given a well defined relation/predicate p , a subject s and a corresponding text about s , we now try to estimate the relation cardinality (i.e., the count of $\langle s, p, * \rangle$ triples), based on cardinality assertions found in the text. We chose four Wikidata predicates that span various domains, *child* (P40), *spouse* (P26), *has part* (P527) of a series of creative works (restricted to novel, book and film series), and *contains administrative territorial entity* (P150). As the text source for subjects of each predicate, we consider sentences containing numbers taken from their respective English Wikipedia articles.

Methodology We approach the problem via sequence labelling, i.e., given a sentence containing at least one number, we aim to determine whether each number in the sentence corresponds to the cardinality of a certain relation. We build a Conditional Random Field (CRF) based model with CRF++ (Kudo, 2005) for each relation, taking as features the context lemmas (window size of 5) around the observed token t , along with bigrams and trigrams containing t .

To generate the training data, we rely on distant supervision, annotating *candidate numbers*² in the text as correct cardinalities whenever they correspond to the exact triple count (count > 0) found in the knowledge base. Otherwise, they are labelled as O (for Others), like the rest of non-number tokens. Table 2 contains for each considered relation (p), the number of subjects ($\#s$) in Wikidata, which have links to English Wikipedia pages and have at least one $\langle s, p, * \rangle$ triple.

We predict the relation cardinality of a given $\langle s, p \rangle$ pair by selecting the number positively annotated with marginal probability—resulting from forward-backward inference—higher than 0.1, and choosing the one with the highest probability if there are several.

Experiments Two experimental settings are considered: *vanilla* refers to the distant supervision approach explained above, while for *only-*

²Numbers that are not labelled as DATE, TIME, DURATION, SET, MONEY and PERCENT by Stanford NE-tagger.

nummod, we only annotate a candidate number as correct cardinality if it modifies a noun, i.e., there is an incoming dependency relation of label *nummod* according to the Stanford Dependency Parser. This is to exclude numbers as in “one of the reasons...” from training examples. We also considered a naive baseline, which chooses a random number from a pool of numbers existing in each text about a certain subject.

Furthermore, to estimate how well KB counts are suited as ground truth, we compare them on the the *child* relation with the manually-created *number of children* (P1971) property from Wikidata.

Evaluation Results We manually annotated the evaluation data with the true relation counts, since the knowledge base is highly incomplete, and thus, the triple counts are often incorrect. Whenever the cardinality matches the true count, we also manually inspected how relevant the textual evidence—the context surrounding the cardinal number—is for the observed relation. Table 2 shows the performance of our CRF-based method in finding the correct relation cardinality, evaluated on manually annotated 20 (*has part*), 100 (*admin. terr. entity*) and 200 (*child* and *spouse*) randomly selected subjects that have at least one object.

The random-number baseline achieves a precision of 5% (*has part*), 3.5% (*admin. territ. entity*), 0% (*spouse*) and 11.2% (*child*). Compared to that, especially using *only-nummod*, our method gives encouraging results for *has part*, *admin. territ. entity* and *child*, with 30-50% precision and around 30% F1-score. For *spouse*, the performance is significantly lower, reasons are discussed below. Furthermore, we can observe that using manual ground truth as training data for the *child* relation can boost performance considerably. Still, the performance is significantly below the state-of-the-art in fact extraction, where *child* triples can be extracted from Wikipedia text with 96% precision (Palomares et al., 2016).

5 Analysis

A qualitative analysis of the training data and evaluation results revealed three aspects that make extracting relation cardinalities difficult.

Quality of Training Data Unlike training data for normal fact extraction, which is generally highly correct (e.g., YAGO claims 95% precision (Suchanek et al., 2007)), taking triple counts

<i>p</i>	# <i>s</i>	<i>baseline</i>	<i>vanilla</i>			<i>only-nummod</i>		
		P	P	R	F1	P	R	F1
has part (creative work series)	261	.050	.333	.316	.324	.353	.316	.333
contains admin. terr. entity	18,000	.034	.390	.188	.254	.548	.200	.293
spouse	45,917	0	.014	.011	.013	.028	.017	.021
child	35,057	.112	.151	.129	.139	.320	.219	.260
child (manual ground truth)	6,408		0.374	0.309	0.338	0.452	0.315	0.371

Table 2: Number of Wikidata entities as subjects (*#s*) of each predicate (*p*), and evaluation results on manually annotated randomly selected subjects that have at least an object.

found in knowledge bases as ground truth generally gives wrong results. For example, our manual evaluation of *child* shows that the triple count from Wikidata is 46% lower than what the texts assert.

As shown by the last row of Table 2, higher quality of training data can considerably boost the performance of cardinality extraction. Unfortunately, manually curated data is generally difficult to obtain. We see two avenues to tackle training data quality:

1. *Filtering ground truth.* Instead of taking the counts of all entities as ground truth, one might trade size for quality, e.g., using popular entities only, as for these there are chances that KBs are more complete.
2. *Incompleteness-resilient distant supervision.* Triple counts in KBs are often lower than what is correct, but rarely too high. Thus, an avenue might be to label all numbers equal or higher than the KB count as correct, instead of only considering the equal ones. Given that different cardinalities could then be labelled as correct, this would require a postprocessing step in which conflicting counts are consolidated.

Compositionality Around 16% of false positives in extracting *child* cardinalities can be attributed to failures in identifying the correct count for, e.g., "They have *two* sons and *one* daughter together; he has *four* children from an earlier relationship." This was also observed for other relations, e.g., "The Qidong county has *4* subdistricts, *17* towns and *3* townships under its jurisdiction." We see two avenues to tackle this problem:

1. *Aggregating numbers.* In training data generation, one could label a sequence of number as correct cardinalities if the sum of the numbers is equal to the relation count. In the prediction step, one might sum up all consecutive cardinalities that are labelled with sufficient confidence.

2. *Learning composition rules.* One may try to learn the composition of counts, for instance, that children are composed of sons and daughters, then try to extract the composing cardinalities.

Linguistic Variance We observe that for the *spouse* relation, expressing the count with cardinal numbers ("He has married *four* times") is only found for 4% of subjects. It is more common to express the count with ordinal numbers, e.g., "John's *first* wife, Mary, ...", which allows us to conclude that the spouse-count for John is at least—and most probably more than—one. An approach to such relations might be to identify ordinal numbers that express lower bounds of relations. Subsequently, one could reason over these bounds and try to infer relation counts.

Our initial motivation was to make sense of the so far ignored large fraction of numbers that express relation cardinalities. However, we noticed quickly that relation cardinalities are frequently also expressed without numbers at all. This is especially true for the case of count zero, which is mostly expressed using negation ("He never married"), and the count one, which is expressed using indefinite articles ("They have a child") or the signal-word *only* ("Their only child, James"). Terms such as *twins* or *trilogy* are also ways to express domain-specific relation cardinalities. We see two avenues to approach this variance:

1. *Translation to numbers.* For the 0's and 1's, a possible approach is to translate certain kinds of negation and indefinite articles into explicit numbers (e.g., "do *not* have *any* children" → "have 0 children").
2. *Word similarity with cardinals.* If a word bears high similarity with cardinal numbers, possibly also in other languages such as Latin or Greek, one might consider it as a candidate number.

6 Conclusion

In this paper we have introduced the problem of relation cardinality extraction. We believe that relation cardinalities can be useful in a variety of tasks. Our next goal is to make distant supervision incompleteness-resilient and to deal with compositionality, hoping that these can improve the precision of our approach. We also aim to take ordinals into account and to experiment with linguistic transformation for the cases of cardinalities 0 and 1, hoping that these could boost the recall.

A limitation of our work is also that we only focus on Wikipedia articles, assume that all statements are about the article’s subject, and just take the statement with the highest confidence. In future work we aim to include a larger article base in combination with named entity recognition, coreference resolution and a truth consolidation step.

Acknowledgments

We thank Werner Nutt and Sebastian Rudolph for their feedback on an earlier version of this work. We thank the anonymous reviewers for their helpful comments. This work has been partially supported by the project “The Call for Recall”, funded by the Free University of Bozen-Bolzano.

References

- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. *DBpedia: A nucleus for a web of open data*. Springer.
- Arun Chaganty and Percy Liang. 2016. How much is 131 million dollars? putting numbers in perspective with compositional descriptions. In *ACL*. pages 578–587.
- Mark Craven and Johan Kumlien. 1999. Constructing biological knowledge bases by extracting information from text sources. In *Proceedings of the Seventh International Conference on Intelligent Systems for Molecular Biology*. pages 77–86.
- Luciano Del Corro and Rainer Gemulla. 2013. ClausIE: clause-based open information extraction. In *WWW*. ACM, pages 355–366.
- Yusra Ibrahim, Mirek Riedewald, and Gerhard Weikum. 2016. Making sense of entities and quantities in web tables. In *CIKM*. pages 1703–1712.
- Taku Kudo. 2005. CRF++: Yet another CRF toolkit. *Software available at <http://crfpp.sourceforge.net>*.
- Xiao Ling and Daniel S Weld. 2010. Temporal information extraction. In *AAAI*. volume 10, pages 1385–1390.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. *ACL (System Demonstrations)* pages 55–60.
- Mausam, Michael Schmitz, Stephen Soderland, Robert Bart, and Oren Etzioni. 2012. Open language learning for information extraction. In *EMNLP*. pages 523–534.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *ACL*. pages 1003–1011.
- Paramita Mirza, Simon Razniewski, and Werner Nutt. 2016. Expanding Wikidatas parenthood information by 178%, or how to mine relation cardinalities. *ISWC Posters & Demos*.
- Tom M. Mitchell, William W. Cohen, Estevam R. Hruschka Jr., Partha Pratim Talukdar, Justin Betteridge, Andrew Carlson, Bhavana Dalvi Mishra, Matthew Gardner, Bryan Kisiel, Jayant Krishnamurthy, Ni Lao, Kathryn Mazaitis, Thahir Mohamed, Ndapandula Nakashole, Emmanouil Antonios Platanios, Alan Ritter, Mehdi Samadi, Burr Settles, Richard C. Wang, Derry Tanti Wijaya, Abhinav Gupta, Xinlei Chen, Abulhair Saparov, Malcolm Greaves, and Joel Welling. 2015. Never-ending learning. In *AAAI*. pages 2302–2310.
- Sebastian Neumaier, Jürgen Umbrich, Josiane Xavier Parreira, and Axel Polleres. 2016. Multi-level semantic labelling of numerical values. In *ISWC*. pages 428–445.
- Thomas Palomares, Youssef Ahres, Juhana Kangaspunta, and Christopher Ré. 2016. Wikipedia knowledge graph with DeepDive. In *ICWSM*. pages 65–71.
- Simon Razniewski, Fabian M. Suchanek, and Werner Nutt. 2016. But what do we actually know? *Proceedings of AKBC* pages 40–44.
- Jannik Strötgen and Michael Gertz. 2010. Heideitime: High quality rule-based extraction and normalization of temporal expressions. In *SemEval Workshop*. pages 321–324.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. YAGO: a core of semantic knowledge. *WWW* pages 697–706.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM* 57(10):78–85.