Extraction of Part-Whole Relations from Turkish Corpora

Tuğba Yıldız, Savaş Yıldırım, and Banu Diri

Department of Computer Engineering, İstanbul Bilgi University
Dolapdere, 34440 İstanbul, Turkey
Department of Computer Engineering, Yıldız Technical University
Davutpaşa, 34349 İstanbul, Turkey
{tdalyan,savasy}@bilgi.edu.tr,
banu@ce.yildiz.edu.tr

Abstract. In this work, we present a model for semi-automatically extracting part-whole relations from a Turkish raw text. The model takes a list of manually prepared seeds to induce syntactic patterns and estimates their reliabilities. It then captures the variations of part-whole candidates from the corpus. To get precise meronymic relationships, the candidates are ranked and selected according to their reliability scores. We use and compare some metrics to evaluate the strength of association between a pattern and matched pairs. We conclude with a discussion of the result and show that the model presented here gives promising results for Turkish text.

Keywords: Meronym, Part-Whole, Semantic Lexicon, Semantic Similarity.

1 Introduction

The meronym has been referred to as a part-whole relation that represents the relationship between a part and its corresponding whole. The discovery of meronym relations plays an important role in many NLP applications, such as question answering, information extraction [1–3], query expansion [4] and formal ontology [5, 6].

Different types of part-whole relations have been proposed in the literature [7–9]. One of the most important and well-known taxonomies, designed by Winston [8] identified part-whole relations as falling into six types: Component-Integral, Member-Collection, Portion-Mass, Stuff-Object, Feature-Activity and Place-Area. On the other hand, the most popular and useful ontologies such as WordNet have also classified meronyms into three types: component-of (HAS-PART), member-of (HAS-MEMBER) and stuff-of (HAS-SUBSTANCE)[10].

A variety of methods have been proposed to identify part-whole relations from a text source. Some studies employed lexico-syntactic patterns for indicating part-whole relations. There have also been other approaches such as statistical, supervised, semi-supervised or WordNet corporation[1, 14, 16–18].

A. Gelbukh (Ed.): CICLing 2013, Part I, LNCS 7816, pp. 126–138, 2013.

This study is a major attempt to semi-automatically extract part-whole relations from a Turkish corpus. Other recent studies to harvest meronym relations and types of meronym relations for Turkish are based on dictionary definition (TDK) and WikiDictionary [11–13].

The rest of this paper is organized as follows: Section 2 presents and compares related works. We explain our methodology in Section 3. Implementation details are explained in Section 4. Experimental results and their evaluation are reported in Section 5.

2 Related Works

Many studies for automatically discovering part-whole relations from text have been based on Hearst's [14] pattern-based approach. Hearst developed a method to identify hyponym (is-a) relation from raw text with using lexico-syntactic patterns. Although the same technique was applied to extract meronym relations in [14], it was reported that efforts concluded without great success.

In [16], it was proposed a statistical methods in very large corpus to find parts. Using Hearst's methods, five lexical patterns and six seeds (book, building, car, hospital, plant, school) for wholes were identified. Part-whole relations extracted by using patterns were ranked according to some statistical criteria with an accuracy of 55% for the top 50 words and an accuracy of 70% for the top 20 words.

A semi-automatic method was presented in [1] for learning semantic constraints to detect part-whole relations. The method picks up pairs from WordNet and searches them on text collection: SemCor and LA Times from TREC-9. Sentences that containing pairs were extracted and manually inspected to obtain list of lexico-syntactic patterns. Training corpus was generated by manually annotating positive and negative examples. C4.5 decision tree was used as learning procedure. The model's accuracy was 83%. The extended version of this study was proposed in [17]. An average precision of 80.95% was obtained.

Hage [3] developed a method to discover part-whole relations from vocabularies and text. The method followed two main phases: learning part-whole patterns and learning wholes by applying the patterns. An average precision of 74% was achieved.

A weakly-supervised algorithm, Espresso [18] used patterns to find several semantic relations besides meronymic relations. The method automatically detected generic patterns to choose correct and incorrect ones and to filter with the reliability scoring of patterns and instances. System performance for part-of relations on TREC was 80% precision.

Another attempt at automatic extraction of part-whole relation was for a Chinese Corpus [19]. The sentence containing part-whole relations was manually picked and then annotated to get lexico-syntactic patterns. Patterns were employed on training corpus to find pairs of concepts. A set of heuristic rules were proposed to confirm part-whole relations. The model performance was evaluated with a precision of 86%.

Another important studies were proposed in [2, 20]. In [2], a set of seeds for each type of part-whole relations was defined. The minimally-supervised information extraction algorithm, Espresso [18] successfully retrieved part-whole relations from corpus. For English corpora, the precision was 80% for general seeds and 82% for structural part-of seeds. In [20], an approach extracted meronym relation from domain-specific text for product development and customer services.

3 Methodology

The pattern-based approach proposed here is implemented in two phases: Pattern identification and part-whole pair detection. Fig. 1 represents how the system is split up into its components and shows data flow among these components. The system takes a huge corpus and a set of unambiguous part-whole pairs. It then proposes a list of parts for a given whole.

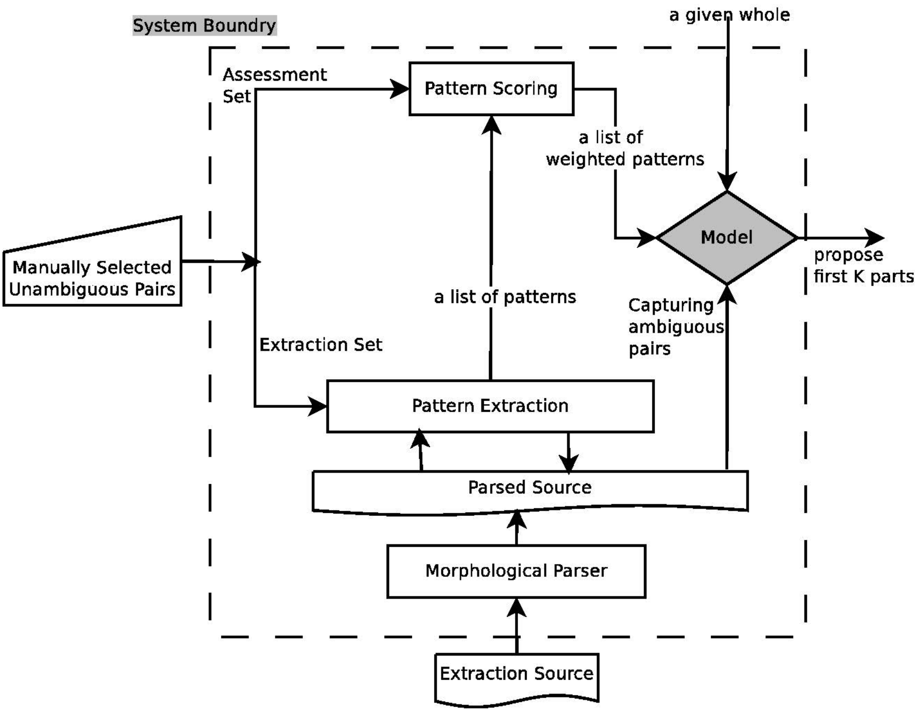

Fig. 1. High-level representation of the system

3.1 Pattern Identification

We begin by manually preparing a set of unambiguous seed pairs that convey a part-whole relation. For instance, the pair (engine, car) would be member of that set. The seed set is further divided into two subsets: an *extraction set* and an *assessment set*.

Each pair in the extraction set is used as query for retrieving sentences containing that pair. Then we generalize many lexico-syntactic expressions by replacing part and whole token with a wildcard or any meta character. The second set, the assessment set, is then used to compute the usefulness or reliability scores of all the generalized patterns. Those patterns whose reliability scores, *rel(p)*, are very low are eliminated. The remaining patterns are kept, along with their reliability scores. A classic way to estimate *rel(p)* of an extraction pattern is to measure how it correctly identifies the parts of a given whole. The success rate is obtained by dividing the number of correctly extracted pairs by the number of all extracted pairs. The outcome of entire phase is a list of reliable lexico-syntactic expressions along with their reliability scores.

3.2 Part-Whole Pair Detection

In order to extract the pairs among which there is a part-whole relation, the previously generated patterns are applied to an extraction source that is a Turkish raw text. The instantiated instances (part-whole pairs) are assessed and ranked according to their reliability scores, where reliability score of a pair is described below.

There are several ways to compute a reliability score for both pattern and instance. In [18], the reliability score of a pattern, *rel(p)*, is proposed as shown in equation (1) and that of an instance, *rel(i)*, is formulated as in equation (2).

$$rel(p) = \frac{\sum_{p \in P} \left(\frac{pmi(i,p)}{max_{pmi}} \times rel(i) \right)}{|P|} \tag{1}$$

$$rel(i) = \frac{\sum_{i \in I} \left(\frac{pmi(i,p)}{max_{pmi}} \times rel(p) \right)}{|I|} \tag{2}$$

pmi is the pointwise mutual information that is one of the commonly used metrics for the strength of association between two variables, where max_{pmi} is the maximum pmi value between all pairs and all patterns and where *rel(i)* is the reliability of instance i. Initially, all reliability scores of instances in set of unambiguous pairs are set to 1. Then, reliability score of a pattern is calculated based on these *rel(i)* scores.

In [18], the pmi score between an instance $i(x, y)$ and pattern p was formulated as in following equation (3).

$$pmi(i,p) = log \frac{|x,p,y|}{|x,*,y||*,p,*|} \tag{3}$$

where $|x, p, y|$ is the number of times instance $i(x, y)$ is instantiated with pattern p, $|x, *, y|$, $|*, p, *|$ are the individual distributions of instance and pattern respectively. However, the defect in the formula is that the pmi score always takes negative values. This leads a ranking the reverse of the expected. It must be multiplied by the numbers of all pairs matched by all patterns, $|*, *, *|$. Thus, we redefined the formula as shown in equation (4).

$$pmi(i, p) = log\frac{|x, p, y||*, *, *|}{|x, *, y||*, p, *|} \tag{4}$$

A frequent pair in a particular pattern does not necessarily convey a part-whole relation. Thus, to calculate reliability of a pair, all patterns are taken into consideration as shown in equation (2).

In our research, we experiment with three different measures of association (pmi, dice, tscore) to evaluate their performance. We also utilized *inverse document frequency (idf)* to cover more specific parts. The motivation for use of *idf* is to differentiate distinctive features from other common ones.

We categorized our parts into two groups; *distinctive* and *general* parts. If a part of a given whole is inheritable from hypernyms of that whole, we call this kind of part general or inheritable. If, not, we call such part specific or distinctive part. E.g. a desk has has-part relationship with drawer and segment as in WordNet. While drawer is distinctive part of desk, segment is a general part that inherits from its hypernym *artifact*. Indeed, it is really difficult to apply this chaining approach to all nouns. Instead of using all hypernym chain, the parts come from *physical entity* or some particular hypernyms and the distinctive ones can be merged.

Thus, to distinguish the distinctiveness, we utilized *idf* that is obtained by dividing the number of times a part occurs in part position by how many pairs retrieved by all the patterns. We observed that the most frequent part instances are top, inside, segment, side, bottom, back, front and state, head etc. All of these resemble general features.

3.3 Baseline Algorithm

Pointwise mutual information and other measures can be alternatively used between part and whole rather than instance pair and pattern. To designate a baseline algorithm, for a given whole, its possible parts are retrieved from a list ranked by association measure between whole and part that are instantiated by a reliable pattern as formulated in equation (5).

$$assoc(whole, part) = \frac{|whole, pattern, part|}{|*, pattern, part||whole, pattern, *|} \tag{5}$$

We intuitively designated a baseline algorithm to compare the results and the expectation is that a proposed model should outperform the baseline algorithm. The baseline function is based on most reliable and productive pattern, the genitive pattern. As Table 1 suggests, the rel(genitive-pattern) has the best score in accordance with average of all three measures (pmi, dice and tscore) and the capacity is about 2M part-whole pairs.

For a given whole, all parts that co-occur with that whole in the genitive pattern are extracted. Taking co-occurrence frequency between the whole and part could be misleading due to some nouns frequently placed in part/head position such as side, front, behind, outside. To overcome the problem, the co-occurrence, the individual distributions of both whole and part must be taken into account as shown in equation (5). These final scores are ranked and their first K parts are selected as the output of baseline algorithm.

4 Experimental Design

In our experiment, we used a set of natural language resources for Turkish; a huge corpus of 500M tokens and a morphological parser provided by [15]. The morphological parser based on a two-level morphology has an accuracy of %98. The web corpus contains four sub-corpora. Three of them are from major Turkish news portals and another corpus is a general sampling of web pages in the Turkish Language. The corpus is tokenized and encoded at paragraph and sentence level.

The morphological parser splits a surface token into its morphemes in system architecture as shown in Figure 1. The representation of a parsed token is in the form of **surface/root/pos/[and all other markers**]. When the genitive phrase "arabanın kapısı (door of the car)" is given, the parser split it into the parts as below.

```
English: (door of the car)
Turkish: arabanın kapısı
Parsed : araba+noun+a3sg+pnon+gen kapı+a3sg+pnon+p3sg
```

In order to identify lexical forms that express part-whole relations, we manually selected 200 seed pairs. Out of 200 pairs, 50 are used as pattern extraction set to extract the lexico-syntactic patterns and 150 are used as assessment set to compute the reliability scores of each pattern, *rel(p)*. All sentences containing *part* and corresponding *whole* token in extraction set are retrieved. Replacing part/whole token with a meta character, e.g. wildcard, we extracted many patterns.

However, due to the noisy nature of the web corpus and the difficulties of an agglutinative language, many patterns have poor extraction capacity. Turkish is a relatively free word order language with agglutinating word structures. The noun phrases can easily change their position in a sentence without changing the meaning of the sentence, and only affecting its emphasis. This is a big challenge for syntactic pattern extraction. Based on reliability scores, we decided to filter out some generated patterns and finally obtained six different significant patterns. Here is the list of the patterns, their examples and related regular expression formula:

1. **Genitive Pattern: NP+gen NP+pos**
   In Turkish, there is only one genitive form: The modifier morphologically takes a genitive case, *Gen (nHn)* and the head takes possessive agreement *pos(sH)* as shown before ("arabanın kapısı/ door of the car"). The morphological feature of genitive is a good indicator to disclose a semantic relation between a head and its modifier. In this case, we found that the genitive has a good indicative capacity, although it can encode various semantic interpretations. Taking the example, *Ali's team*, the first interpretation could be that *the team belongs to Ali*, the second interpretation is that *Ali's favorite team* or *the team he supports*. To overcome such problem, researcher have done many studies based on statistical evidence, some well-known semantic similarity measurements and semantic constraints based on world knowledge resources.
   The regular expression of genitive pattern for "'arabanın kapısı' is as follows:

```
Regex  : \w+\+noun[\w\+]+gen +\w+\+noun[\w\+]+p3sg
```

2. **NP+nom NP+pos**

```
English: (car door)
Turkish: araba kapısı
Parsed : araba+noun+a3sg+pnon+nom  kapı+a3sg+pnon+p3sg
Regex  : \w+\+noun\+a3sg\+pnon\+nom \w+\+noun\+[\w\+]+p3sg
```

3. **NP+Gen (N—ADJ)+ NP+Pos**

```
English: (back garden gate of the house)
Turkish: Evin arka bahçe kapısı
Parsed : Evin+ev+noun+a3sg+pnon+gen
         arka+arka+noun+a3sg+pnon+nom
         bahçe+bahçe+noun+a3sg+pnon+nom
         kapısı+kapı+noun+a3sg+p3sg+nom

Regex: \w+\+noun[\w\+]+gen
       (\w+\+noun\+a3sg\+pnon\+nom |\w+\+adj[\w\+]+ )
       \w+\+noun[\w\+]+p3sg
```

4. **NP of one-of NPs**

```
English: (the door of one of the houses)
Turkish: Evlerden birinin kapısı
Parsed : Evlerden+ev+noun+a3pl+pnon+abl
         birinin+biri+pron+quant+a3sg+p3sg+gen
         kapısı+kapı+noun+a3sg+p3sg+nom
Regex  : \w+\+noun\+a3pl\+pnon\+abl
         birinin\+biri\+pron\+quant\+a3sg\+p3sg\+gen
         \w+\+noun\+\w+\+p3sg
```

5. **NP whose NP**

```
English: The house whose door is locked
Turkish: Kapısı kilitli olan ev
Parsed:  Kapısı+kapı+noun+a3sg+p3sg+nom
         kilitli+kilit+noun+a3sg+pnon+nom-adj*with
         olan+ol+verb+pos-adj*prespart
         ev+ev+noun+a3sg+pnon+nom
Regex :  \w+\+noun[\w\+]+p3sg\+\w+
         (\w+\+noun\+a3sg\+pnon\+nom|
          \w+\+noun\+a3sg\+pnon\+nom\-adj\*with)
         (\w+\+verb\+pos\-adj\*prespart|
          \w+\+verb\+pos\+narr\+a3sg) \w+\+noun\+a3sg
```

6. **NP with NPs**

```
English: the house with garden and pool
Turkish: bahçeli ve havuzlu ev
Parsed : bahçeli+bahçeli+adj ve+ve+conj
         havuzlu+havuz+noun+a3sg+pnon+nom-adj*with
         ev+ev+noun+a3sg+pnon+nom
Regex:   \w+\+noun\+a3sg\+pnon\+nom\-adj\*with \w+\+noun\+
```

All patterns are evaluated according to their usefulness. To assess them, output of each pattern is checked against a given assessment set. Setting instance reliability of all pairs in the set to 1, reliability score of the patterns are computed as shown before. For a assessment set size of 150 pairs, all pattern and their *rel(p)* are given in Table 1

When comparing the patterns, P1 is the most reliable pattern with respect to all measures. P1 is based on genitive case which many studies utilized it for the problem. We roughly order the pattern as P1, P2, P3, P6, P4, P5 by their normalized average scores in the Table 1.

Table 1. Reliability of Patterns

	rel(P1)	**rel(P2)**	**rel(P3)**	**rel(P4)**	**rel(P5)**	**rel(P6)**
pmi	1.58	1.53	0.45	0.04	0.07	0.57
dice	0.01	0.003	0.01	0.004	0.001	0.003
tscore	0.11	0.12	0.022	0.0004	0.001	0.03

To calculate reliability of instances, we utilize not only pmi measure, but also dice, t-score and idf measures. In equation(1), *rel(p)* and equation (2), *rel(i)*, association measure can be **pmi, pmi-idf, dice, dice-idf, tscore,** and **tscore-idf**. For a particular whole noun, all possible parts instantiated by patterns are selected as a candidate set. For each association measure, their *rel(p)* and *rel(i)* scores are calculated and further sorted. The first K candidate parts are checked against the expected parts.

5 Evaluation

For the evaluation phase, we manually and randomly selected five whole words: *book, computer, ship, gun* and *building.* For each whole noun, the experimental results are given in Table 2.

Table 2. The results of the scores for five wholes

whole	**pmi**	**pmi-idf**	**dice**	**dice-idf**	**tscore**	**tscore-idf**	**baseline**	**average**
gun-10	2	**4**	1	1	0	1	2	1.57
gun-20	4	**5**	3	2	1	1	4	2.86
gun-30	**6**	**6**	**6**	**6**	2	3	**6**	5
book-10	9	3	**10**	**10**	8	7	8	7.86
book-20	**18**	9	**18**	**18**	16	12	13	14.86
book-30	22	14	22	**23**	21	20	17	19.86
building-10	4	2	5	**7**	**7**	6	**7**	5.43
building-20	11	8	**15**	14	**15**	13	**15**	13
building-30	17	13	22	**23**	20	19	18	18.86
ship-10	**9**	7	**9**	**9**	6	5	**9**	7.71
ship-20	14	13	**18**	**18**	9	10	15	13.86
ship-30	18	17	**26**	24	13	14	21	19
computer-10	8	**9**	**9**	**9**	6	7	8	8
computer-20	**16**	15	13	15	8	11	10	12.57
computer-30	**21**	16	20	20	10	15	14	16.57
average prec.								
precision10	64%	50%	68%	**72%**	54%	52%	68%	61.14%
precision20	63%	50%	**67%**	**67%**	49%	47%	57%	57.14%
precision30	56%	44%	**64%**	**64%**	44%	47.3%	50.6%	52.86%

Where gun-10 means that we evaluated first 10 selection of all measures for whole gun. For a better evaluation, we selected first 10, 20 and 30 candidates ranked by the association measure defined above. The proposed parts were manually evaluated by looking at their semantic role. We needed to differentiate part-whole relations from other possible meanings. Indeed, all the proposed parts are somehow strongly associated with corresponding whole. However, our specific goal here is to discover meronymic relationship and, thus we tested our results with respect to the component-integral meronymic relationship as defined in [8] or HAS-PART in WordNet. For first 10 output, dice-idf with precision of 72% performs better than others on average. For first 20 selection, dice and dice-idf share the highest scores of 67%. For first 30 selection, dice, dice-idf with precision of 64% outperforms other measures.

Looking at the Table 2, for the first 10 selection, all measures perform well against all wholes but gun. This is simply because *gun* gives less corpus evidence to discover parts of it. With a deeper observation, we have manually captured only 9 distinctive parts and 10 general parts, whereas whole *building* has 51 parts, out of which 13 are general parts.

We conducted another experiment to distinguish distinctive parts from general ones. Excluding general parts from the expected list, we re-evaluated the result of the experiments. The results were, of course, less successful but a better fine-grained model was obtained. The result are shown in Table 3. The table showed that all idf weighted measures are better than others. For the first 30 selection, when idf is applied, pmi, dice measures are increased by 2% and tscore measure is increased by 7.3% on average as expected.

Table 3. The results for distinctive parts

precision	pmi	pmi-idf	dice	dice-idf	tscore	tscore-idf	baseline	average
precision10	50	50	58	**64**	34	44	60	51.43
precision20	48	50	48	**53**	34	40	51	46.29
precision30	40.67	42.67	47.33	**49.33**	31.33	38.67	40.67	41.52

General parts can easily captured when running the system for *entity* or any hypernym. To do so, we checked noun *"şey"*‘ (thing) to cover more general parts or features. We retrieved some meaningful nouns such as *top, end, side, base, front, inside, back, out* as well as other meaningless parts.

Additionally, we can easily apply is-a relation, whereas we cannot always apply the same principle to part-whole hierarchy. For instance if tail is a meronym of cat and tiger is a hyponym of cat, by inheritance, tail must be a meronym of tiger then. However, transitivity could be limited in the part-whole relation. Handle is meronym of door, door is a meronym of house. It can incorrectly implied that the house has a handle. On the other hand, finger-hand-body hierarchy is a workable example to say that a body has a finger.

As our another result, Table 4 partly confirms our expectation that the success rate from a larger training seed set is slightly better than those from a smaller one. As we increase the seed size from 50 to 150, only pmi measure clearly improves and the other measures did now show significant improvements.

Table 4. The precision (prec) results for training set (TS) size of 50,100 and 150

#of_TS	results	pmi	pmi-idf	dice	dice-idf	tscore	tscore-idf	baseline	avg.
train50	prec10	64	52	**70**	68	52	44	68	59.71
	prec20	59	50	70	**68**	48	45	57	56.71
	prec30	51.33	43.33	62.67	**64**	42	46	50.67	51.43
train100	prec10	68	50	**72**	70	52	48	66	60.86
	prec20	63	49	**68**	66	48	44	58	56.57
	prec30	56	44.67	63.33	**64**	42.67	45.33	50.67	52.38
train150	prec10	64	50	68	**72**	54	52	68	61.14
	prec20	63	50	**67**	**67**	49	47	57	57.14
	prec30	56	44	**64**	**64**	44	47.33	50.67	52.86

The goal of the study is to retrieve meronymic relation, more specifically component integral (CI) relation. Looking at the result in Table 5, almost those candidates that are incorrectly proposed in terms of component integral relations, however, fall into other semantic relations such as property, cause etc. When evaluating the results with respect to these semantic relations SR that includes CI as well, we obtain better precision but coarser relations. The average score for first 10 selection is 68 and 90 with respect to CI and SR respectively. For first 20 and 30 selection, we can conclude that system successfully disclose semantic relation in coarser manner.

The last remark is that only dice and dice-idf significantly outperformed the baseline algorithm. Looking at Table 2 and Table 3, all other measures did not show significant performance with respect to the baseline algorithm.

Table 5. The results for Component-Integral (CI) relation and other Semantic Relations(SRs)

whole	**#CI**	**#Other SRs**
gun-10	1	5
gun-20	3	11
gun-30	6	14
building-10	5	10
building-20	15	19
building-30	22	28
computer-10	9	10
computer-20	13	16
computer-30	20	24
ship-10	9	10
ship-20	18	20
ship-30	26	28
book-10	10	10
book-20	18	18
book-30	22	25
average-10	68	90
average-20	67	84
average-30	64	79.3

6 Conclusion

In this study, we proposed a model that semi-automatically acquires meronymy relations from a Turkish raw text. The study is a major corpus driven attempt to extract part-whole relations from a Turkish corpus. The raw text has been tokenized and morphologically parsed. Some manually prepared seeds were used to induce lexico-syntactic patterns and determine their usefulness. Six reliable patterns were extracted and scored. Based on these patterns, the model captures a variety of part-whole candidates from the parsed corpus.

We conducted several experiments on a huge corpus of size of 500M tokens. All experiments indicate that proposed model has promising results for Turkish Language. According to experimental results and observation, we conclude some points as follows.

We compared the strength of some association measure with respect to their precisions. Looking at the first 10 selection of each measure, all measures are equally same in terms of precision, whereas for first 30 selection dice and dice-idf outperformed other ones with precision of 64%. Looking at the all resulting tables, we can say that only dice and dice-idf significantly outperformed the baseline algorithm. The performance of other measures were not confirmed our expectation.

For distinctive part retrieval rather than general parts, we expected that idf weighted measures are significantly better than others. For the first 30 selection, when idf is applied, pmi, dice measures are increased by 2% and tscore measure is increased by 7.3% on average as expected. To get general part for a given whole, hypernymy and meronymy relation can be intertwined. While we cannot always apply part-whole hierarchy due to lack of transitivity, hypernymy relation can be easily and safely applied. We showed that some common parts can be hierarchically inherited from artifact or entity synset.

We conducted an experiment to see the importance of larger seed set. It partly confirms our expectation that the success rate from a larger training seed set is slightly better than those from a smaller one. As we increase the seed size from 50 to 150, only pmi measure clearly improves and the other measures did now show significant improvements.

Initial goal of the study is to retrieve meronymic relation, more specifically component-integral relation. Most of the incorrectly retrieved candidates actually fall into other kind of semantic relations such as has-a, property, cause etc. When evaluating the results with respect to these semantic relations, we obtained better precision but coarser relations as expected.

References

1. Girju, R., Badulescu, A., Moldovan, D.I.: Learning Semantic Constraints for the Automatic Discovery of Part-Whole Relations. In: Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL 2003), pp. 1–8. Association for Computational Linguistics, Morristown (2003)
2. Ittoo, A., Bouma, G.: On Learning Subtypes of the Part-Whole Relation: Do Not Mix Your Seeds. In: Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, pp. 1328–1336 (2010)
3. Van Hage, W.R., Kolb, H., Schreiber, G.: A Method for Learning Part-Whole Relations. In: Cruz, I., Decker, S., Allemang, D., Preist, C., Schwabe, D., Mika, P., Uschold, M., Aroyo, L.M. (eds.) ISWC 2006. LNCS, vol. 4273, pp. 723–735. Springer, Heidelberg (2006)

4. Buscaldi, D., Rosso, P., Arnal, E.S.: Using the WordNet Ontology in the GeoCLEF Geographical Information Retrieval Task. In: Peters, C., Gey, F.C., Gonzalo, J., Müller, H., Jones, G.J.F., Kluck, M., Magnini, B., de Rijke, M., Giampiccolo, D. (eds.) CLEF 2005. LNCS, vol. 4022, pp. 939–946. Springer, Heidelberg (2006)
5. Artale, A., Franconi, E., Guarino, N., Pazzi, L.: Part-Whole Relations in Object-Centered Systems: An Overview. Data Knowledge Engineering 20(3), 347–383 (1996)
6. Schulz, S., Hahn, U.: Part-whole representation and reasoning in formal biomedical ontologies. Artificial Intelligence in Medicine 34(3), 179–200 (2005)
7. Madelyn, A.I.: Problems of the part-whole relation. In: Evens, M. (ed.) Relational Models of the Lexicon, pp. 261–288. Cambridge University Press, New York (1989)
8. Winston, M.E., Chaffin, R., Herrmann, D.: A Taxonomy of Part-Whole Relations. Cognitive Science 11(4), 417–444 (1987)
9. Keet, C.M., Artale, A.: Representing and reasoning over a taxonomy of part-whole relations. Applied Ontology 3(1-2), 91–110 (2008)
10. Miller, G.A., Beckwith, R., Fellbaum, C., Gross, D., Miller, K.: WordNet: An on-line lexical database. International Journal of Lexicography 3, 235–244 (1990)
11. Yazıcı, E., Amasyalı, M.F.: Automatic Extraction of Semantic Relationships Using Turkish Dictionary Definitions. EMO Bilimsel Dergi 1(1), 1–13 (2011)
12. Ayşe Şerbetçi, A., Orhan, Z., Pehlivan, İ.: Extraction of Semantic Word Relations in Turkish from Dictionary Definitions. In: Proceedings of the ACL 2011 Workshop on Relational Models of Semantics (RELMS 2011), Portland, Oregon, USA, pp. 11–18. Association for Computational Linguistics (2011)
13. Orhan, Z., Pehlivan, İ., Uslan, V., Önder, P.: Automated Extraction of Semantic Word Relations in Turkish Lexicon. Mathematical and Computational Applications 16(1), 13–22 (2011)
14. Hearst, M.A.: Automatic Acquisition of Hyponyms from Large Text Corpora. In: COLING 1992 Proceedings of the 14th conference on Computational linguistics, vol. 2, pp. 539–545. Association for Computational Linguistics (1992)
15. Sak, H., Güngör, T., Saraçlar, M.: Turkish Language Resources: Morphological Parser, Morphological Disambiguator and Web Corpus. In: Nordström, B., Ranta, A. (eds.) GoTAL 2008. LNCS (LNAI), vol. 5221, pp. 417–427. Springer, Heidelberg (2008)
16. Berland, M., Charniak, E.: Finding Parts in Very Large Corpora. In: Proceedings of the 37th Annual meeting of the Association for Computational Linguistics on Computational Linguistics, pp. 57–64 (1999)
17. Girju, R., Badulescu, A., Moldovan, D.I.: Automatic Discovery of Part-Whole Relations. Computational Linguistics 32(1), 83–135 (2006)
18. Pantel, P., Pennacchiotti, M.: Espresso: Leveraging Generic Patterns for Automatically Harvesting Semantic Relations. In: Proceedings of the 21st International Conference on Computational Linguistics and the 44th Annual Meeting of the ACL, pp. 113–120. ACL Press, Sydney (2006)
19. Cao, X., Cao, C., Wang, S., Lu, H.: Extracting Part-Whole Relations from Unstructured Chinese Corpus. In: Proceedings of the 2008 Fifth International Conference on Fuzzy Systems and Knowledge Discovery, pp. 175–179. IEEE Computer Society (2008)
20. Ittoo, A., Bouma, G., Maruster, L., Wortmann, H.: Extracting Meronymy Relationships from Domain-Specific, Textual Corporate Databases. In: Hopfe, C.J., Rezgui, Y., Métais, E., Preece, A., Li, H. (eds.) NLDB 2010. LNCS, vol. 6177, pp. 48–59. Springer, Heidelberg (2010)