

USAAR at SemEval-2016 Task 13: Hyponym Endocentricity

Liling Tan¹, Francis Bond² and Josef van Genabith¹
Universität des Saarlandes¹ / Campus, Saarbrücken, Germany
Nanyang Technological University² / 14 Nanyang Drive, Singapore
liling.tan@uni-saarland.de, bond@ieee.org
josef.van_genabith@dfki.de

Abstract

This paper describes our submission to the SemEval-2016 Taxonomy Extraction Evaluation (TExEval-2) Task. We examine the endocentric nature of hyponyms and propose a simple rule-based method to identify hypernyms at high precision. For the food domain, we extract lists of terms from the Wikipedia lists of lists by using the name of each list as the endocentric head and treating all terms in the extracted tables as the hyponym of the endocentric head.

Our submission achieved competitive results in taxonomy construction and ranked top in hypernym identification when evaluated against gold standard taxonomies and also in manual evaluation of novel relations not covered by the gold standard taxonomies.

1 Introduction

Semantic taxonomies provide structured world knowledge to Artificial Intelligence (AI) and Natural Language Processing (NLP) systems. Traditional broad-coverage taxonomies such as CYC (Lenat, 1995), SUMO (Pease et al., 2002; Miller, 1995), YAGO (Suchanek et al., 2007) and Freebase (Bollacker et al., 2008) have been manually created or curated with much effort.

With the rapid technological evolution, it is more feasible to construct a domain-specific taxonomy that caters to domain or company specific terminology (Lefever, 2015). This motivated the move towards unsupervised approaches to taxonomy extraction (Berland and Charniak, 1999; Lin and Pantel,

2001; Snow et al., 2006) and specifically focused towards particular domains (Velardi et al., 2013; Bordea et al., 2015).

The aim of the Taxonomy Extraction Evaluation (TExEval) task is to automatically find lexical relations between pairs of terms within several specified domains. Previously, we have developed a hypernym extraction system using word embeddings by exploiting the frequent occurrence of the ‘*X is a Y*’ pattern in encyclopedic text (Tan et al., 2015).

We have achieved competitive results in SemEval-2015 and as a follow up to our study, we would like to explore the endocentric nature of hyponyms that contributed substantially to the system performance in the previous TExEval task.

Below, we will briefly (i) describe related work on different approaches to taxonomy induction, (ii) explain the linguistic phenomenon of endocentricity, (iii) present our endocentric hypo-hypernym identification system and the results of our submission to the TExEval-2 task in SemEval-2016.

2 Related Work

The hierarchical structure of domain concepts is made of hypo-hypernymy relations between terms. Different approaches have been proposed to induce these relations automatically ranging from pattern/rule-based approaches (Hearst, 1992; Girju, 2003; Kozareva et al., 2008; Ceesay and Hou, 2015) to clustering and frequency based approaches (Lin, 1998; Caraballo, 2001; Pantel and Ravichandran, 2004; Grefenstette, 2015), classification approaches (Snow et al., 2004; Ritter et al., 2009; Espinosa Anke et al., 2015) and graph-based ap-

proaches (Kozareva and Hovy, 2010; Navigli et al., 2011; Fountain and Lapata, 2012; Tuan et al., 2014; Cleuziou et al., 2015).

More recently, there is a resurgence of vector space or distributional approaches (Van Der Plas, 2005; Lenci and Benotto, 2012; Santus et al., 2014) primarily because of the renaissance of deep learning and neural networks.

Semantic knowledge can be thought of as a vector space where each word is presented by a point and the proximity between words in this space quantifies their semantic association. The vector space is usually constructed from the distribution of words across contexts such that similar meanings tend to be found close to each other within the vector space (Mitchell and Lapata, 2010).

With the present advancement in neural nets and word embeddings (Mikolov et al., 2013; Pennington et al., 2014; Levy et al., 2014; Shazeer et al., 2016), neural space models are gaining popularity in taxonomy induction and relation extraction tasks (Saxe et al., 2013; Fu et al., 2014; Tan et al., 2015).

3 Endocentricity

Early research in theoretical linguistics discussed the idea of *endocentric* vs. *exocentric* constructions (Brugmann, 1886; Aleksandrov, 1886; Brockelmann, 1908; Bloomfield, 1983).

A grammatical construction is *endocentric* when it fulfills the same linguistic function as one of its part(s). For instance, the word *goldfish* is an endocentric compound noun as syntactically it functions as a noun just as its component part *fish* and semantically the compound denotes a type of *fish*.

Conversely, when a grammatical construction made of two or more parts is *exocentric*, no part component carries the linguistic function or meaning assigned to the complex construction. Intuitively, we would expect that there are many endocentric hyponyms in a taxonomy where part of the term conveys its main meaning and usually that part of term would be its hypernym.

The endo/exocentricity feature of a lexical term assumes that the term can be split into two or more parts. For example, *fish* is a single noun that cannot be split, thus it cannot be endo- or exocentric.

While experimenting with ways to weight a term

for information retrieval, Jones (1979) observed that compound nouns often follow the head-modifier principle where the meaning of the term can be conveyed by part(s) of the compound. Approaching endocentricity from a different angle, Nichols et al. (2005) identified the semantic head(s) of a term as its hypernym using the lowest scoping element of the Robust Minimal Recursion Semantics (RMRS) (Copestake et al., 2005) structure of the dictionary definition of the term.

In the first TExEval task in SemEval-2015, both Lefever (2015) and Tan et al. (2015)¹ independently developed string-based systems that exploit the endocentric nature of hyponyms.

In our submission to the TExEval-2 task (Bordea et al., 2016), we seek to answer the question of exactly *“how many hyponyms within a taxonomy are endocentric?”*. Additionally, we exploit the endocentric nature of the hyponyms to extend the taxonomy by crawling and cleaning Wikipedia’s *List of Lists of Lists*.² Often these lists of terms are found in Wikipedia marked up as tables or in bullet forms.

4 Identifying Endocentric Parts

The main implementation of the rule-based identifier³ checks *if a term T1 is a substring of another term T2 and if so, assign T1 as a hypernym of T2*. Examples of hypo-hypernym pairs captured by this rule includes are (*psycholinguistics, linguistics*), (*kobe beef, beef*), (*sauce gribiche, sauce*).

Our implementation is simpler than the three part morpho-syntactic analyzer component of the multimodal taxonomy constructor in Lefever (2015). She implemented rules for three different syntactic constructions which check for suffixes and treat single-word terms and multi-word terms differently while our implementation is agnostic to the single and multi-word distinction.

In addition to the first rule, *if a term contains the “of” preposition, we swap the assignment and check that T2 starts with T1 then assign T2 as a hypernym of T1*. Examples of hypo-hypernym pairs

¹<https://github.com/alvations/USAAR-SemEval-2015/tree/master/task17-USAAR-WLV>

²https://en.wikipedia.org/wiki/List_of_lists_of_lists

³Our open-source implementation can be found at <https://github.com/alvations/Endocentricity>

captured by this swap rule are (*elixir of life, elixir*), (*sociology of education, sociology*).

To improve the precision of the identifier, we set a threshold of a minimum character length of three when identifying a term as a hypernym.

5 Extending a Taxonomy with Wikipedia List of Lists of Lists

The Wikipedia List of Lists of Lists (LOLOL) is a crowdsourced list of lists of terms. We adapted the customized crawler⁴ (Tan et al., 2014; Tan and Ordan, 2015) to crawl for tables or bullet points in the Wikipedia subpages of the LOLOL for the food domain. We started the crawl from these seed pages under the bullet point of https://en.wikipedia.org/wiki/List_of_lists_of_lists#Food_and_drink.

When the crawler lands on each List of Lists (LOL) page, it will *treat the URL suffix as the hypernym* and *find words in the bullet points or tables that contain endocentric hyponyms*.

If an endocentric hyponym exists, it will extract either (i) all the terms in bold font if the LOL page is bulleted or (ii) all terms in the first column if the LOL page is in table form. The choice of the first column is based on the fact that often LOL tables are bi-column, one containing the terms and the other the gloss or/and description of the term.

5.1 Limitations of LOLOL Trawler

There are a number of issues with this *trawling* (*crawl+clean*) approach to extend the taxonomy.

LOL pages are not standardized: The way the crawler cleans the bullets or tables on each LOL page is not standardized because there is no constraint put on the format of the Wikipedia’s LOL page. Our trawler only managed to crawl and clean less than 30 LOL pages when extracting the new terms for the food domain.

LOL pages are inceptive: The depth of how nested the LOLs are is undefined. Our trawler can start with a `List_of_foods` page and it leads to the `List_of_breads` page and then the `List_of_American_breads` page and it contin-

⁴It was built for crawling translations and diachronic texts in previous SemEval tasks

ues. For sanity, we had to break our trawler at the second page depth and return to the main LOLOL page to move on to the next LOL that we have not previously trawled.

6 Results

Table 1 presents the overview results of our submissions to the TExEval-2 task. Only the results for the food domain contains the hypo-hypernym pairs extracted by our trawler. The rest of the domains comprise of the outputs solely generated by our endocentric hypo-hypernym identifier.

Although it is counter-intuitive to think that endocentric hypo-hypernym pairs can be wrong, this example aptly demonstrates the limitation: (*honey bunches of oats, honey*). In this case, neither ‘*honey bunches of oats*’ can be a hypernym of ‘*honey*’ nor vice versa.

When compared against the gold standard taxonomies, our submission achieved the highest precision in the environment, food (WordNet), science (Eurovoc) and science (WordNet) domains.

As for the Food domain, we expected the fall in precision due to the additional terms that we introduced from the Wikipedia LOLOL outside of the gold standard taxonomy. Thus, we are also unable to determine the true “correctness” of these terms (indicated by the dash in Table 1).

Looking at the proportion of the number of hypo-hypernym pairs that our system correctly identified, we can approximate that *15-25% of the hypernyms in a taxonomy can be easily identified through their endocentric hyponyms* by taking the ratio of $\#Correct / \#Terms$.

However, the proportions presented in Table 1 exclude the correct hypo-hypernym pairs that are identified but are not currently in the gold-standard taxonomy. Table 2 presents the results of the manual evaluation for the precision of 100 randomly selected hypo-hypernym pairs that are not in the gold standard taxonomy. Our system achieved top precision in all domains other than science (Eurovoc).

If we consider the precision scores from Table 2 as the precision of the remaining identified but not correct hypo-hypernym pairs in Table 1, we would be able to add to the 15-25% hypononym endocentricity in taxonomies. However, the aggregation of

	Environment (Eurovoc)	Food (WordNet)	Food	Science (Eurovoc)	Science (WordNet)	Science
#Terms	261	1486	1555	370	125	452
#Relations	261	1576	1587	452	124	465
#Correct / Identified	38 / 47	381 / 540	- / 4347*	66 / 104	25 / 30	119 / 312
Precision	0.8085	0.7056	0.0603	0.6333	0.8173	0.3814
Recall	0.1456	0.2418	0.1651	0.1532	0.1881	0.2559
F-score	0.2468	0.3601	0.0883	0.2468	0.3058	0.3063
F&M	0.0007	0.0021	0.0	0.0023	0.0008	0.0020

Table 1: Results of our Endocentric Hypo-Hypernym Identifier Against the Gold Standard Taxonomy (#Terms refers to the no. of terms in the domain and #Relations refers to the no. of hypo-hypernym pairs found in the gold-standard taxonomy. #Correct / #Identified refers to the proportion of hypo-hypernym pairs our system has correctly identified. **Bold** items indicates that it is highest score among the participating teams in TExEval-2. The asterisk * indicates that the trawler was used to produce submissions for this domain.)

Domain	JUNLP	TAXI	NUIG-UNLP	USAAR	QASSIT
Environment (Eurovoc)	0.02	0.11	0.08	0.22	0.07
Food	0.20	0.36	-	0.73*	-
Food (Wordnet)	0.18	0.32	-	0.81	-
Science	0.06	0.14	0.09	0.71	0.07
Science (Eurovoc)	0.02	0.02	0.04	0.00	0.05
Science (Wordnet)	0.06	0.22	0.05	0.47	0.22

Table 2: Results of Manual Evaluation on 100 Random Novel Hypo-Hypernym Pairs for Participating Teams In TExEval-2

the manual evaluation results should only be considered if the novel hypo-hypernym relations are curated and added to the standard taxonomies.

		Baseline	TAXI	USAAR
Food	P	0.5000	0.3388	0.7056
(WordNet)	R	0.2576	0.2932	0.2418
Science	P	0.6897	0.3747	0.8173
(WordNet)	R	0.2655	0.3805	0.1881

Table 3: Comparison of String-based Methods

Comparing against the TExEval-2 organizers baseline string-based method and the TAXI lexico-syntactic substring approach (Panchenko and Biemann, 2016) for the WordNet taxonomies, our system achieved highest precision but underperformed in recall as shown in Table 3.

Since our main implementation of our hypernym identifier is language independent, in retrospect, we can easily remove the swap rule that is attached to the English ‘of’ and apply the identifier to other languages in the TExEval-2 task.

6.1 Other Participating Systems

Table 2 presents a summary of the results of novel hypo-hypernym pairs identified by the participating systems in TExEval-2. A detailed overview of the results of TExEval-2 is presented in Bordea et al. (2016).

JUNLP relied on substrings and relations extracted from BabelNet (Navigli and Ponzetto, 2012) to identify hyper-hyponym pairs. Although it is sensible to approach the task using an existing ontology, their system achieved relatively low precision on the manual evaluation of novel hyper-hyponym pairs. The NUIG-UNLP team extended previous work on vector space approaches to taxonomy induction by comparing the similarity between the dense word embeddings of the hyponyms and their candidate hypernyms. They system achieved high recall but attained low precision (Pocostales, 2016).

Similar to our endocentric-based approach, the TAXI team extended the substring-based approach by filtering the hypernym candidates based on corpora statistics of lexico-syntactic patterns. Additionally, they applied pruning methods to improve the

ontological structure which resulted in high Fowlkes and Mallows (F&M) Measure (Panchenko and Biemann, 2016). QASSIT used lexical patterns to extract hypernym candidates and applied the pretopological space graph optimization technique that is based on genetic algorithm to achieve the desired taxonomy structure (Cleuziou and Moreno, 2016).

TAXI and QASSIT ranked first and second in the taxonomy construction criterion of the TExEval task. Both teams used graph pruning techniques to improve the taxonomy structure and implicitly improve the F&M scores⁵ of their taxonomy. Although our endocentricity based hypo-hypernym extraction system ranked first in hypernym identification of TExEval task, we ranked third in taxonomy construction with an overall F&M score of 0.0013.

7 Conclusion

In this paper, we have described our submission to the Taxonomy Extraction Evaluation (TExEval-2) Task for SemEval-2016. We have empirically shown that 15-25% of the hypernyms in a taxonomy can be easily identified through their endocentric hyponyms and we briefly discuss the intuitions and limitations of the approach.

We have achieved competitive results in taxonomy construction and achieved top precision score for hypernym identification in most domains involved in the task.

Acknowledgments

The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007-2013/ under REA grant agreement no 317471.

References

- Alexander Aleksandrov. 1886. *Sprachliches aus dem Nationaldichter Litauens Donalitus. T. 1, Zur Semasiologie : Inaugural-Dissertation zur Erlangung des Grades eines Magisters der vergleichenden Sprachkunde*. Schnakenburg.
- Matthew Berland and Eugene Charniak. 1999. Finding Parts in Very Large Corpora. In *Proceedings of the*

⁵Overall F&M scores for TAXI and QASSIT are 0.4064 and 0.2908

37th annual meeting of the Association for Computational Linguistics on Computational Linguistics, pages 57–64.

- Leonard Bloomfield. 1983. *An Introduction to the Study of Language*. 2. Benjamins.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 1247–1250, New York, NY, USA. ACM.
- Georgeta Bordea, Paul Buitelaar, Stefano Faralli, and Roberto Navigli. 2015. Semeval-2015 task 17: Taxonomy extraction evaluation (texeval). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 902–910, Denver, Colorado.
- Georgeta Bordea, Els Lefever, and Paul Buitelaar. 2016. Semeval-2016 task 13: Taxonomy extraction evaluation (texeval-2). In *Proceedings of the 10th International Workshop on Semantic Evaluation*.
- Carl Brockelmann. 1908. *Grundriss der vergleichenden Grammatik der semitischen Sprachen*, volume 1. Reuther & Reichard.
- Karl Brugmann. 1886. *Vergleichende Grammatik der indogermanischen Sprachen*. Walter de Gruyter.
- Sharon Ann Caraballo. 2001. *Automatic Construction of a Hypernym-labeled Noun Hierarchy from Text*. Ph.D. thesis, Providence, RI, USA. AAI3006696.
- Bamfa Ceesay and Wen Juan Hou. 2015. Ntnu: An unsupervised knowledge approach for taxonomy extraction. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 938–943, Denver, Colorado.
- Guillaume Cleuziou and Jose G. Moreno. 2016. Qassit at semeval-2016 task 13: On the integration of semantic vectors in pretopological spaces for lexical taxonomy acquisition. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval 2016)*, San Diego, California. Association for Computational Linguistics.
- Guillaume Cleuziou, Davide Buscaldi, Gaël Dias, Vincent Levorato, and Christine Largeton. 2015. Qassit: A pretopological framework for the automatic construction of lexical taxonomies from raw texts. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 955–959, Denver, Colorado.
- Ann Copestake, Dan Flickinger, Carl Pollard, and Ivan A Sag. 2005. Minimal recursion semantics: An introduction. *Research on Language and Computation*, 3(2-3):281–332.

- Luis Espinosa Anke, Horacio Saggion, and Francesco Ronzano. 2015. Taln-upf: Taxonomy learning exploiting crf-based hypernym extraction on encyclopedic definitions. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 949–954, Denver, Colorado.
- Trevor Fountain and Mirella Lapata. 2012. Taxonomy Induction using Hierarchical Random Graphs. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 466–476.
- Ruiji Fu, Jiang Guo, Bing Qin, Wanxiang Che, Haifeng Wang, and Ting Liu. 2014. Learning Semantic Hierarchies via Word Embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1199–1209.
- Roxana Girju. 2003. Automatic Detection of Causal Relations for Question Answering. In *Proceedings of the ACL 2003 workshop on Multilingual summarization and question answering-Volume 12*, pages 76–83.
- Gregory Grefenstette. 2015. Inriasac: Simple hypernym extraction methods. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 911–914, Denver, Colorado.
- Marti A Hearst. 1992. Automatic Acquisition of Hyponyms from Large Text Corpora. In *Proceedings of the 14th conference on Computational linguistics-Volume 2*, pages 539–545.
- Karen Sparck Jones. 1979. Experiments in relevance weighting of search terms. *Information Processing & Management*, 15(3):133–144.
- Zornitsa Kozareva and Eduard Hovy. 2010. A Semi-Supervised Method to Learn and Construct Taxonomies using the Web. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1110–1118.
- Zornitsa Kozareva, Ellen Riloff, and Eduard Hovy. 2008. Semantic Class Learning from the Web with Hyponym Pattern Linkage Graphs. In *Proceedings of ACL-08: HLT*, pages 1048–1056, Columbus, Ohio.
- Els Lefever. 2015. Lt3: A multi-modular approach to automatic taxonomy construction. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 944–948, Denver, Colorado.
- Douglas B Lenat. 1995. CYC: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11):33–38.
- Alessandro Lenci and Giulia Benotto. 2012. Identifying hypernyms in distributional semantic spaces. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*, pages 75–79.
- Omer Levy, Yoav Goldberg, and Israel Ramat-Gan. 2014. Linguistic regularities in sparse and explicit word representations. In *CoNLL*, pages 171–180.
- Dekang Lin and Patrick Pantel. 2001. Discovery of Inference Rules for Question-Answering. *Natural Language Engineering*, 7(04):343–360.
- Dekang Lin. 1998. Automatic Retrieval and Clustering of Similar Words. In *Proceedings of the 17th international conference on Computational linguistics-Volume 2*, pages 768–774.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- George A. Miller. 1995. WordNet: A Lexical Database for English. *Commun. ACM*, 38(11):39–41.
- Jeff Mitchell and Mirella Lapata. 2010. Composition in Distributional Models of Semantics. *Cognitive Science*, 34(8):1388–1439.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. Babelnet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*, 193:217–250.
- Roberto Navigli, Paola Velardi, and Stefano Faralli. 2011. A Graph-Based Algorithm for Inducing Lexical Taxonomies from Scratch. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 1872–1877.
- Eric Nichols, Francis Bond, and Dan Flickinger. 2005. Robust ontology acquisition from machine-readable dictionaries. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 1111–1116.
- Stefano Ruppert Eugen Remus Steffen Naets Hubert Fairon Cedrick Ponzetto Simone Paolo Panchenko, Alexander Faralli and Chris Biemann. 2016. Taxi at semeval-2016 task 13: a taxonomy induction method based on lexico-syntactic patterns, substrings and focused crawling. In *Proceedings of the 10th International Workshop on Semantic Evaluation*. Association for Computational Linguistics.
- Patrick Pantel and Deepak Ravichandran. 2004. Automatically Labeling Semantic Classes. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*.
- Adam Pease, Ian Niles, and John Li. 2002. The Suggested Upper Merged Ontology: A Large Ontology for

- the Semantic Web and its Applications. In *In Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web*.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543.
- Joel Pocostales. 2016. Nuig-unlp at semeval-2016 task 13: A simple word embedding-based approach for taxonomy extraction. In *Proceedings of the 10th International Workshop on Semantic Evaluation*.
- Alan Ritter, Stephen Soderland, and Oren Etzioni. 2009. What is this, anyway: Automatic hypernym discovery. In *AAAI Spring Symposium: Learning by Reading and Learning to Read*, pages 88–93.
- Enrico Santus, Alessandro Lenci, Qin Lu, and Sabine Schulte im Walde. 2014. Chasing hypernyms in vector spaces with entropy. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 38–42, Gothenburg, Sweden.
- Andrew M. Saxe, James L. McClelland, and Surya Ganguli. 2013. Learning Hierarchical Category Structure in Deep Neural Networks. pages 1271–1276.
- Noam Shazeer, Ryan Doherty, Colin Evans, and Chris Waterson. 2016. Swivel: Improving embeddings by noticing what’s missing. *arXiv preprint arXiv:1602.02215*.
- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2004. Learning syntactic patterns for automatic hypernym discovery. *Advances in Neural Information Processing Systems 17*.
- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2006. Semantic Taxonomy Induction from Heterogenous Evidence. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pages 801–808.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: A Core of Semantic Knowledge. In *Proceedings of the 16th International Conference on World Wide Web, WWW ’07*, pages 697–706, New York, NY, USA. ACM.
- Liling Tan and Noam Ordan. 2015. Usaar-chronos: Crawling the web for temporal annotations. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 846–850, Denver, Colorado.
- Liling Tan, Anne Schumann, Jose Martinez, and Francis Bond. 2014. Sensible: L2 translation assistance by emulating the manual post-editing process. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 541–545, Dublin, Ireland.
- Liling Tan, Rohit Gupta, and Josef van Genabith. 2015. Usaar-wlv: Hypernym generation with deep neural nets. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 932–937, Denver, Colorado.
- Luu Anh Tuan, Jung-jae Kim, and Kiong See Ng. 2014. Taxonomy construction using syntactic contextual evidence. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 810–819.
- Lonneke Van Der Plas. 2005. Automatic acquisition of lexico-semantic knowledge for qa. In *Proceedings of the IJCNLP workshop on Ontologies and Lexical Resources*, pages 76–84.
- Paola Velardi, Stefano Faralli, and Roberto Navigli. 2013. OntoLearn Reloaded: A Graph-based Algorithm for Taxonomy Induction. *Computational Linguistics*, 39(3):665–707.