

CogALex-V Shared Task: LexNET - Integrated Path-based and Distributional Method for the Identification of Semantic Relations

Vered Shwartz

Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel
vered1986@gmail.com

Ido Dagan

Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel
dagan@cs.biu.ac.il

Abstract

We present a submission to the CogALex 2016 shared task on the corpus-based identification of semantic relations, using LexNET (Shwartz and Dagan, 2016), an integrated path-based and distributional method for semantic relation classification. The reported results in the shared task bring this submission to the third place on subtask 1 (word relatedness), and the first place on subtask 2 (semantic relation classification), demonstrating the utility of integrating the complementary path-based and distributional information sources in recognizing concrete semantic relations. Combined with a common similarity measure, LexNET performs fairly good on the word relatedness task (subtask 1). The relatively low performance of LexNET and all other systems on subtask 2, however, confirms the difficulty of the semantic relation classification task, and stresses the need to develop additional methods for this task.

1 Introduction

Discovering whether words are semantically related and identifying the specific semantic relation that holds between them is a key component in many NLP applications, such as question answering and recognizing textual entailment (Dagan et al., 2013). Automated methods for semantic relation identification are commonly corpus-based, and mainly rely on the distributional representation of each word.

The CogALex shared task on the corpus-based identification of semantic relations consists of two subtasks. In the first task, the system needs to identify for a word pair whether the words are semantically related or not (e.g. *True:(dog, cat)*, *False:(dog, fruit)*). In the second task, the goal is to determine the specific semantic relation that holds for a given pair, if any (*PART_OF:(tail, cat)*, *HYPER:(cat, animal)*).

In this paper we describe our approach and system setup for the shared task. We use LexNET (Shwartz and Dagan, 2016), an integrated path-based and distributional method for semantic relation classification. LexNET was the system with the overall best performance on subtask 2, and was ranked third on subtask 1, demonstrating the utility of integrating the complementary path-based and distributional information sources in recognizing semantic relatedness.¹

To aid in recognizing whether a pair of words are related at all (subtask 1), we combine LexNET with a common similarity measure (cosine similarity), achieving fairly good performance, and a slight improvement upon using cosine similarity alone. Subtask 2, however, has shown to be extremely difficult, with LexNET and all other systems achieving relatively low F_1 scores. The conflict between the mediocre performance and the recent success of distributional methods on several other common datasets for semantic relation classification (Baroni et al., 2012; Weeds et al., 2014; Roller et al., 2014) could be explained by the stricter evaluation setup in this subtask, which is supposed to demonstrate more closely real-world application settings. The difficulty of the semantic relation classification task emphasizes the need to develop better methods for this task.

This work is licenced under a Creative Commons Attribution 4.0 International License. License details:
<http://creativecommons.org/licenses/by/4.0/>

¹LexNET's code is available at <https://github.com/vered1986/LexNET>, and the shared task results are available at <https://sites.google.com/site/cogalex2016/home/shared-task/results>

2 Background

2.1 Word Relatedness

Recognizing word relatedness is typically addressed by distributional methods. To determine to what extent two terms x and y are related, a vector similarity or distance measure is applied to their distributional representations: $sim(\vec{v}_{w_x}, \vec{v}_{w_y})$. This is a straightforward application of the distributional hypothesis (Harris, 1954), according to which related words occur in similar contexts, hence have similar vector representations.

Most commonly, vector cosine is adopted as a similarity measure (Turney et al., 2010). Many other measures exist, including but not limited to Euclidean distance, KL divergence (Cover and Thomas, 2012), Jaccard’s coefficient (Salton and McGill, 1986), and more recently neighbor rank (Hare et al., 2009; Lapesa and Evert, 2013) and APSyn (Santus et al., 2016a).² To turn this task into a binary classification task, where x and y are classified as either related or not, one can set a threshold to separate similarity scores of related and unrelated word pairs.

2.2 Semantic Relation Classification

Recognizing lexical semantic relations between words is valuable for many NLP applications, such as ontology learning, question answering, and recognizing textual entailment. Most corpus-based methods classify the relation between a pair of words x and y based on the distributional representation of each word (Baroni et al., 2012; Roller et al., 2014; Fu et al., 2014; Weeds et al., 2014). Earlier methods utilized the dependency paths that connect the joint occurrences of x and y in the corpus as a cue to the relation between the words (A. Hearst, 1992; Snow et al., 2004; Nakashole et al., 2012). Recently, Shwartz and Dagan (2016) presented LexNET, an extension of HypeNET (Shwartz et al., 2016). This method integrates both path-based and distributional information for semantic relation classification, which outperformed approaches that rely on a single information source, on several common datasets (Baroni and Lenci, 2011; Neculescu et al., 2015; Santus et al., 2015; Santus et al., 2016b).

3 System Description

In LexNET, a word-pair (x, y) is represented as a feature vector, consisting of a concatenation of distributional and path-based features: $\vec{v}_{xy} = [\vec{v}_{w_x}, \vec{v}_{paths(x,y)}, \vec{v}_{w_y}]$, where \vec{v}_{w_x} and \vec{v}_{w_y} are x and y ’s word embeddings, providing their distributional representation, and $\vec{v}_{paths(x,y)}$ is the average embedding vector of all the dependency paths that connect x and y in the corpus. Dependency paths are embedded using a LSTM (Hochreiter and Schmidhuber, 1997), as described in Shwartz et al. (2016). This vector is then fed into a neural network that outputs the class distribution \vec{c} , and then the pair is classified to the relation with the highest score r :

$$\vec{c} = \text{softmax}(\text{MLP}(\vec{v}_{xy})) \quad (1a)$$

$$r = \text{argmax}_i \vec{c}[i] \quad (1b)$$

MLP stands for Multi Layer Perceptron, and could be computed with or without a hidden layer (equations 2 and 3, respectively):

$$\vec{h} = \tanh(W_1 \cdot \vec{v}_{xy} + b_1) \quad (2a)$$

$$\text{MLP}(\vec{v}_{xy}) = W_2 \cdot \vec{h} + b_2 \quad (2b)$$

$$\text{MLP}(\vec{v}_{xy}) = W_1 \cdot \vec{v}_{xy} + b_1 \quad (3)$$

where W_i and b_i are the network parameters and \vec{h} is the hidden layer.

²See Lee (1999) for an extensive list of such measures.

	Method	Hyper-parameters	Corpus size	P	R	F ₁
Subtask 1	Cos	word2vec, $t: 0.3$	100B	0.759	0.795	0.776
	LexNET	hidden layers: 0, dropout: 0.0, epochs: 3	6B	0.780	0.561	0.652
	LexNET+Cos	word2vec, $w_L = 0.3, w_C = 0.7, t = 0.29$	~100B	0.814	0.854	0.833
Subtask 2	Dist	dep-based, method: concat, classifier: SVM, L_1	3B	0.611	0.598	0.600
	LexNET	hidden layers: 0, dropout: 0.0, epochs: 5	6B	0.658	0.646	0.642

Table 1: Performance scores on the validation set along with hyper-parameters and effective corpus size (#tokens) used by each method. Subtask 2 results refer to the subset of related pairs, as detailed in § 4.2.

3.1 A Note About Word Relatedness

While path-based approaches have been commonly used for semantic relation classification (A. Hearst, 1992; Snow et al., 2004; Nakashole et al., 2012; Necsulescu et al., 2015), they have never been used for word relatedness, which is considered a “classical” task for distributional methods. We argue that path-based information can improve performance of word relatedness tasks as well (see Section 4.1). We train LexNET to distinguish between two classes: RELATED and UNRELATED, and combine it with the common cosine similarity measure to tackle subtask 1.

4 Experimental Settings

The shared task organizers provided a dataset extracted from EVALution 1.0 (Santus et al., 2015), which was split into training and test sets. As instructed, we trained and tuned our method on the training set, and evaluated it once on the test set. To tune the hyper-parameters, we split the training set to 90% train and 10% validation sets. Since the dataset contains only 318 different words in the x slot, we performed the split such that the train and the validation contain distinct x words.³

LexNET has several tunable hyper-parameters. Similarly to Shwartz and Dagan (2016), we used the English Wikipedia dump from May 2015 as an underlying corpus (3B tokens), and initialized the network’s word embeddings with the 50-dimensional pre-trained GloVe word embeddings (Pennington et al., 2014), trained on Wikipedia and Gigaword 5 (6B tokens). We fixed this hyper-parameter due to computational limitations with higher-dimensional embeddings. For each subtask, we tuned LexNET’s hyper-parameters on the validation set: the number of hidden layers (0 or 1), the number of training epochs, and the word dropout rate (see Shwartz et al. (2016) for technical details). Table 1 displays the best performing hyper-parameters in each subtask, along with the performance on the validation set, which is detailed below.

4.1 Subtask 1: Word Relatedness

We tuned LexNET’s hyper-parameters on the validation set, disregarding the similarity measure at this point, and then chose the model that performed best on the validation set and combined it with the similarity measure.

We computed cosine similarity for each (x, y) pair in the dataset: $\cos(\vec{v}_{w_x}, \vec{v}_{w_y}) = \frac{\vec{v}_{w_x} \cdot \vec{v}_{w_y}}{\|\vec{v}_{w_x}\| \cdot \|\vec{v}_{w_y}\|}$, and normalized it to the range $[0, 1]$. We scored each (x, y) pair by a combination of LexNET’s score for the RELATED class and the cosine similarity score:

$$\text{Rel}(x, y) = w_C \cdot \cos(\vec{v}_{w_x}, \vec{v}_{w_y}) + w_L \cdot \bar{c}[\text{RELATED}] \quad (4)$$

where w_C, w_L are the weights assigned to cosine similarity and LexNET’s scores respectively, such that $w_C + w_L = 1$. We tuned the weights and a threshold t using the validation set, and classified (x, y) as related if $\text{Rel}(x, y) \geq t$. The word vectors used to compute the cosine similarity scores were chosen among several available pre-trained embeddings.⁴ For completeness we also report the performance of two baselines: cosine similarity ($w_C = 1$) and LexNET ($w_L = 1$, fixed $t = 0.5$).

³A random split yielded perfect results on the validation set, which were due to lexical memorization (Levy et al., 2015).

⁴word2vec (300 dimensions, SGNS, trained on GoogleNews, 100B tokens) (Mikolov et al., 2013), GloVe (50-300 dimensions, trained on Wikipedia and Gigaword 5, 6B tokens) (Pennington et al., 2014), and dependency-based embeddings (300 dimensions, trained on Wikipedia, 3B tokens) (Levy and Goldberg, 2014).

	Method	P	R	F_1
Subtask 1	Random Baseline	0.283	0.503	0.362
	Majority Baseline	0.000	0.000	0.000
	Cos	0.841	0.672	0.747
	LexNET+Cos	0.754	0.777	0.765
Subtask 2	Random Baseline	0.073	0.201	0.106
	Majority Baseline	0.000	0.000	0.000
	Dist	0.469	0.371	0.411
	LexNET	0.480	0.418	0.445

Table 2: Performance scores on the test set in each subtask, of the selected methods and the baselines.

4.2 Subtask 2: Semantic Relation Classification

The subtask’s train set is highly imbalanced towards random instances (roughly 10 times more than any other relation), and training any supervised method leads to overfitting to the random class. We therefore trained the model only on the related classes (excluding RANDOM pairs), for which the classes are more balanced. During inference time, we used the model from subtask 1 to assign a relatedness score to each pair, $Rel(x, y)$, and computed the class distribution using the model from subtask 2, only for pairs that were related according to this score.

Finally, we applied a heuristic that if for a word pair (x, y) , the difference in scores between the top scoring classes is low (< 0.2), and the top class is SYN, then it is only classified as SYN if the number of paths between x and y is smaller than 3. This is due to the fact that synonyms are hard to recognize with both distributional and path-based approaches (Shwartz and Dagan, 2016), but it is known that they do not tend to co-occur.

To compare LexNET’s performance on the validation set with other methods’ performances, we adapted the distributional baseline employed by Shwartz et al. (2016) and Shwartz and Dagan (2016), where a classifier is trained on the combination of x and y ’s word embeddings. We experimented with several combination methods (concatenation (Baroni et al., 2012), difference (Fu et al., 2014; Weeds et al., 2014), and ASYM (Roller et al., 2014)), regularization factors, and pre-trained word embeddings (Mikolov et al., 2013; Pennington et al., 2014; Levy and Goldberg, 2014). This time, we used cosine similarity (subtask 1) to separate related from unrelated pairs, and trained the classifier only to distinguish between the related classes. Similarly to subtask 1, we tune LexNET and the baseline’s hyper-parameters on the validation set. The best performance is reported in Table 1.

5 Results and Analysis

Table 2 displays the performance of our methods and the baselines on the test set. In addition to the two baselines provided by the shared task organizers (majority and random), we report also the results of our baselines detailed in Section 4. The majority baseline classifies all the instances as UNREALTED (subtask 1) or RANDOM (subtask 2). Since these labels are excluded from the averaged F_1 computation, this baseline’s performance metrics are all zero.

Subtask 1: Word Relatedness Cos achieves fairly good performance ($F_1 = 0.747$), and LexNET+Cos slightly improves upon it. To better understand LexNET’s contribution, we examined pairs that were correctly classified by LexNET+Cos while being incorrectly classified by Cos. Out of the 57 pairs that were true negative in LexNET and false positive in Cos, we judged only one as somehow related (*(death, man)*).

We sampled 25 (from the 184) true positive pairs in LexNET+Cos that were false negatives in Cos, and found that they were all connected via paths in the corpus, suggesting that LexNET’s contribution comes also from the path-based component, rather than only from adding distributional information. 12 of the pairs contained a polysemous term, for which the relation holds in a specific sense (e.g. *(fire, shoot)*). 5 other pairs had a weak relation, e.g. *(compact, car)*. While a *car* can be *compact*, non of these words is one of the most related words to the other.⁵ As noted by Shwartz and Dagan (2016), these are

⁵*car* is mostly related to *driver*, *cars*, and *race*, and *compact* to *compactness* and *locally*.

		predictions								
		ANT	RANDOM	HYPER	PART_OF	SYN	relation	P	R	F_1
gold	ANT	40.28	30.28	5.56	5.83	18.06	ANT	0.450	0.403	0.425
	RANDOM	2.35	93.07	1.31	1.37	1.90	RANDOM	0.897	0.931	0.914
	HYPER	10.21	22.77	45.81	6.02	15.18	HYPER	0.616	0.458	0.526
	PART_OF	2.23	37.05	6.70	47.77	6.25	PART_OF	0.510	0.478	0.493
	SYN	25.96	20.43	14.47	7.23	31.91	SYN	0.278	0.319	0.297

Figure 1: Left: confusion matrix of `LexNET`'s predictions to the subtask 2 test set. Rows indicate gold labels and columns indicate predictions. The value in $[i, j]$ is the percentage of pairs classified to relation j of those labeled i . Right: Per-relation F_1 scores of `LexNET`'s predictions to the test set of subtask 2.

cases in which distributional methods may fail to identify the relation between the words, while even a single meaningful path connecting x and y can capture the relation between them.

Subtask 2: Semantic Relation Classification We note that the overall results on this task are low, in contrast to the success of several methods on common datasets (Baroni et al., 2012; Weeds et al., 2014; Roller et al., 2014; Shwartz and Dagan, 2016). One possible explanation is the stricter and more informative evaluation, that considers the `RANDOM` class as noise, discarding it from the F_1 average.⁶ Additionally, the dataset is lexically split, disabling lexical memorization (Levy et al., 2015). However, the strict evaluation spots a light on the difficulty of this task, which was somewhat obfuscated by the strong results published so far, but might have been obtained thanks to dataset and evaluation peculiarities (Levy et al., 2015; Santus et al., 2016b; Shwartz and Dagan, 2016).

Figure 1 displays `LexNET`'s per relation F_1 scores on the test set, with the corresponding confusion matrix. While the F_1 scores of individual classes are relatively low, the confusion matrix shows that pairs were always classified to the correct relation more than to any other class. A common error comes from subtask 1's model: while most unrelated pairs were classified as unrelated, many related pairs were also classified as unrelated. This may be solved in the future by learning the two subtasks jointly rather than applying a pipeline.

Among the other relations, the performance on synonyms was the worst. The path-based component is weak in recognizing synonyms, which do not tend to co-occur. The distributional information causes confusion between synonyms and antonyms, since both tend to occur in the same contexts. Moreover, synonyms were also sometimes mistaken with hypernyms, as the difference between the two relations is often subtle (Shwartz et al., 2016).

6 Conclusion

We have presented our submission to the CogALex 2016 shared task on corpus-based identification of semantic relations. The submission is based on `LexNET` (Shwartz and Dagan, 2016), an integrated path-based and distributional method for semantic relation classification. `LexNET` was the best-performing system on subtask 2, demonstrating the utility of integrating the complementary path-based and distributional information sources in recognizing semantic relatedness.

We have shown that subtask 1 (word relatedness) reaches reasonable performance with cosine similarity, and is slightly improved when combined with `LexNET`, especially when the relation between the words is non-prototypical. The performance on subtask 2, however, was relatively low for all systems that participated in the shared task, including `LexNET`. This demonstrates the difficulty of the semantic relation classification task, and emphasizes the need to develop improved methods for this task, possibly using additional sources of information.

Acknowledgments

This work was partially supported by an Intel ICRI-CI grant, the Israel Science Foundation grant 880/12, and the German Research Foundation through the German-Israeli Project Cooperation (DIP, grant DA 1600/1-1).

⁶When the random class is included in the averaged F_1 score, the results are: P = 0.780, R = 0.786, F_1 = 0.781.

References

- Marti A. Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *COLING 1992*.
- Marco Baroni and Alessandro Lenci. 2011. Proceedings of the gems 2011 workshop on geometrical models of natural language semantics. pages 1–10.
- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *Proceedings of EACL 2012*, pages 23–32.
- Thomas M Cover and Joy A Thomas. 2012. *Elements of information theory*. John Wiley & Sons.
- Ido Dagan, Dan Roth, and Mark Sammons. 2013. Recognizing textual entailment.
- Ruiji Fu, Jiang Guo, Bing Qin, Wanxiang Che, Haifeng Wang, and Ting Liu. 2014. Learning semantic hierarchies via word embeddings. In *Proceedings of ACL 2014*, pages 1199–1209.
- Mary Hare, Michael Jones, Caroline Thomson, Sarah Kelly, and Ken McRae. 2009. Activating event knowledge. *Cognition*.
- Zellig S Harris. 1954. Distributional structure. *Word*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*.
- Gabriella Lapesa and Stefan Evert. 2013. Evaluating neighbor rank and distance measures as predictors of semantic priming. In *ACL Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2013)*.
- Lillian Lee. 1999. Measures of distributional similarity. In *Proceedings of ACL 1999*.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of ACL 2014*, pages 302–308.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015. Do supervised distributional methods really learn lexical inference relations? In *Proceedings of NAACL-HLT 2015*, pages 970–976.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S Corrado, and Jeffrey Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119.
- Ndapandula Nakashole, Gerhard Weikum, and Fabian Suchanek. 2012. Patty: A taxonomy of relational patterns with semantic types. In *Proceedings of the 2012 Joint Conference EMNLP and CoNLL*, pages 1135–1145.
- Silvia Necsulescu, Sara Mendes, David Jurgens, Núria Bel, and Roberto Navigli. 2015. Reading between the lines: Overcoming data sparsity for accurate classification of lexical relationships. In *Proceedings of *SEM 2015*, pages 182–192.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of EMNLP 2014*, pages 1532–1543.
- Stephen Roller, Katrin Erk, and Gemma Boleda. 2014. Inclusive yet selective: Supervised distributional hypernymy detection. In *Proceedings of COLING 2014*, pages 1025–1036.
- Gerard Salton and Michael J McGill. 1986. Introduction to modern information retrieval.
- Enrico Santus, Frances Yung, Alessandro Lenci, and Chu-Ren Huang. 2015. Proceedings of the 4th workshop on linked data in linguistics: Resources and applications. pages 64–69.
- Enrico Santus, Emmanuele Chersoni, Alessandro Lenci, Chu-Ren Huang, and Philippe Blache. 2016a. Testing apsyn against vector cosine on similarity estimation. *PACLIC*.
- Enrico Santus, Alessandro Lenci, Tin-Shing Chiu, Qin Lu, and Chu-Ren Huang. 2016b. Nine features in a random forest to learn taxonomical semantic relations. In *LREC*.
- Vered Shwartz and Ido Dagan. 2016. Path-based vs. distributional information in recognizing lexical semantic relations. *Proceedings of the 5th Workshop on Cognitive Aspects of the Lexicon (CogALex-V)*.
- Vered Shwartz, Yoav Goldberg, and Ido Dagan. 2016. Improving hypernymy detection with an integrated path-based and distributional method. In *Proceedings of ACL 2016*, pages 2389–2398.
- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2004. Learning syntactic patterns for automatic hypernym discovery. In *NIPS*.
- Peter D Turney, Patrick Pantel, et al. 2010. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37(1):141–188.
- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. Learning to distinguish hypernyms and co-hyponyms. In *Proceedings of COLING 2014*, pages 2249–2259.