

Improving Hypernymy Detection with an Integrated Path-based and Distributional Method

Vered Shwartz Yoav Goldberg Ido Dagan

Computer Science Department

Bar-Ilan University

Ramat-Gan, Israel

vered1986@gmail.com yoav.goldberg@gmail.com dagan@cs.biu.ac.il

Abstract

Detecting hypernymy relations is a key task in NLP, which is addressed in the literature using two complementary approaches. Distributional methods, whose supervised variants are the current best performers, and path-based methods, which received less research attention. We suggest an improved path-based algorithm, in which the dependency paths are encoded using a recurrent neural network, that achieves results comparable to distributional methods. We then extend the approach to integrate both path-based and distributional signals, significantly improving upon the state-of-the-art on this task.

1 Introduction

Hypernymy is an important lexical-semantic relation for NLP tasks. For instance, knowing that *Tom Cruise* is an *actor* can help a question answering system answer the question “which actors are involved in Scientology?”. While semantic taxonomies, like WordNet (Fellbaum, 1998), define hypernymy relations between word types, they are limited in scope and domain. Therefore, automated methods have been developed to determine, for a given term-pair (x, y) , whether y is an hypernym of x , based on their occurrences in a large corpus.

For a couple of decades, this task has been addressed by two types of approaches: distributional, and path-based. In distributional methods, the decision whether y is a hypernym of x is based on the distributional representations of these terms. Lately, with the popularity of word embeddings (Mikolov et al., 2013), most focus has shifted towards supervised distributional methods, in which

each (x, y) term-pair is represented using some combination of the terms’ embedding vectors.

In contrast to distributional methods, in which the decision is based on the *separate* contexts of x and y , path-based methods base the decision on the lexico-syntactic paths connecting the *joint* occurrences of x and y in a corpus. Hearst (1992) identified a small set of frequent paths that indicate hypernymy, e.g. *Y such as X*. Snow et al. (2004) represented each (x, y) term-pair as the multiset of dependency paths connecting their co-occurrences in a corpus, and trained a classifier to predict hypernymy, based on these features.

Using individual paths as features results in a huge, sparse feature space. While some paths are rare, they often consist of certain unimportant components. For instance, “Spelt is a species of wheat” and “Fantasy is a genre of fiction” yield two different paths: *X be species of Y* and *X be genre of Y*, while both indicating that X is-a Y . A possible solution is to generalize paths by replacing words along the path with their part-of-speech tags or with wild cards, as done in the PATTY system (Nakashole et al., 2012).

Overall, the state-of-the-art path-based methods perform worse than the distributional ones. This stems from a major limitation of path-based methods: they require that the terms of the pair occur together in the corpus, limiting the recall of these methods. While distributional methods have no such requirement, they are usually less precise in detecting a specific semantic relation like hypernymy, and perform best on detecting broad semantic similarity between terms. Though these approaches seem complementary, there has been rather little work on integrating them (Mirkin et al., 2006; Kaji and Kitsuregawa, 2008).

In this paper, we present HypeNET, an integrated path-based and distributional method for hypernymy detection. Inspired by recent progress

in relation classification, we use a long short-term memory (LSTM) network (Hochreiter and Schmidhuber, 1997) to encode dependency paths. In order to create enough training data for our network, we followed previous methodology of constructing a dataset based on knowledge resources.

We first show that our path-based approach, on its own, substantially improves performance over prior path-based methods, yielding performance comparable to state-of-the-art distributional methods. Our analysis suggests that the neural path representation enables better generalizations. While coarse-grained generalizations, such as replacing a word by its POS tag, capture mostly syntactic similarities between paths, HypeNET captures also semantic similarities.

We then show that we can easily integrate distributional signals in the network. The integration results confirm that the distributional and path-based signals indeed provide complementary information, with the combined model yielding an improvement of up to 14 F_1 points over each individual model.¹

2 Background

We introduce the two main approaches for hypernymy detection: distributional (Section 2.1), and path-based (Section 2.2). We then discuss the recent use of recurrent neural networks in the related task of relation classification (Section 2.3).

2.1 Distributional Methods

Hypernymy detection is commonly addressed using distributional methods. In these methods, the decision whether y is a hypernym of x is based on the distributional representations of the two terms, i.e., the contexts with which each term occurs *separately* in the corpus.

Earlier methods developed unsupervised measures for hypernymy, starting with symmetric similarity measures (Lin, 1998), and followed by directional measures based on the distributional inclusion hypothesis (Weeds and Weir, 2003; Kotlerman et al., 2010). This hypothesis states that the contexts of a hyponym are expected to be largely included in those of its hypernym. More recent work (Santus et al., 2014; Rimell, 2014) introduce new measures, based on the assumption that the

most typical linguistic contexts of a hypernym are less informative than those of its hyponyms.

More recently, the focus of the distributional approach shifted to supervised methods. In these methods, the (x, y) term-pair is represented by a feature vector, and a classifier is trained on these vectors to predict hypernymy. Several methods are used to represent term-pairs as a combination of each term’s embeddings vector: concatenation $\vec{x} \oplus \vec{y}$ (Baroni et al., 2012), difference $\vec{y} - \vec{x}$ (Roller et al., 2014; Weeds et al., 2014), and dot-product $\vec{x} \cdot \vec{y}$. Using neural word embeddings (Mikolov et al., 2013; Pennington et al., 2014), these methods are easy to apply, and show good results (Baroni et al., 2012; Roller et al., 2014).

2.2 Path-based Methods

A different approach to detecting hypernymy between a pair of terms (x, y) considers the lexico-syntactic paths that connect the *joint* occurrences of x and y in a large corpus. Automatic acquisition of hypernyms from free text, based on such paths, was first proposed by Hearst (1992), who identified a small set of lexico-syntactic paths that indicate hypernymy relations (e.g. *Y such as X, X and other Y*).

In a later work, Snow et al. (2004) learned to detect hypernymy. Rather than searching for specific paths that indicate hypernymy, they represent each (x, y) term-pair as the multiset of all dependency paths that connect x and y in the corpus, and train a logistic regression classifier to predict whether y is a hypernym of x , based on these paths.

Paths that indicate hypernymy are those that were assigned high weights by the classifier. The paths identified by this method were shown to subsume those found by Hearst (1992), yielding improved performance. Variations of Snow et al.’s (2004) method were later used in tasks such as taxonomy construction (Snow et al., 2006; Kozareva and Hovy, 2010; Carlson et al., 2010; Riedel et al., 2013), analogy identification (Turney, 2006), and definition extraction (Borg et al., 2009; Navigli and Velardi, 2010).

A major limitation in relying on lexico-syntactic paths is the sparsity of the feature space. Since similar paths may somewhat vary at the lexical level, generalizing such variations into more abstract paths can increase recall. The PATTY algorithm (Nakashole et al., 2012) applied such generalizations for the purpose of acquiring a taxon-

¹Our code and data are available in:
<https://github.com/vered1986/HypeNET>




Figure 1: An example dependency tree of the sentence “parrot is a bird”, with $x=parrot$ and $y=bird$, represented in our notation as $X/NOUN/nsubj/< be/VERB/ROOT/-Y/NOUN/attr/>$.

omy of term relations from free text. For each path, they added generalized versions in which a subset of words along the path were replaced by either their POS tags, their ontological types or wild-cards. This generalization increased recall while maintaining the same level of precision.

2.3 RNNs for Relation Classification

Relation classification is a related task whose goal is to classify the relation that is expressed between two target terms in a given sentence to one of predefined relation classes. To illustrate, consider the following sentence, from the SemEval-2010 relation classification task dataset (Hendrickx et al., 2009): “The [apples] $_{e_1}$ are in the [basket] $_{e_2}$ ”. Here, the relation expressed between the target entities is *Content – Container*(e_1, e_2).

The shortest dependency paths between the target entities were shown to be informative for this task (Fundel et al., 2007). Recently, deep learning techniques showed good performance in capturing the indicative information in such paths.

In particular, several papers show improved performance using recurrent neural networks (RNN) that process a dependency path edge-by-edge. Xu et al. (2015; 2016) apply a separate long short-term memory (LSTM) network to each sequence of words, POS tags, dependency labels and WordNet hypernyms along the path. A max-pooling layer on the LSTM outputs is used as the input of a network that predicts the classification. Other papers suggest incorporating additional network architectures to further improve performance (Nguyen and Grishman, 2015; Liu et al., 2015).

While relation classification and hypernymy detection are both concerned with identifying semantic relations that hold for pairs of terms, they differ in a major respect. In relation classification the relation should be expressed in the given text, while in hypernymy detection, the goal is to recognize a generic lexical-semantic relation between terms that holds in many contexts. Accordingly, in relation classification a term-pair is represented

by a single dependency path, while in hypernymy detection it is represented by the multiset of all dependency paths in which they co-occur in the corpus.

3 LSTM-based Hypernymy Detection

We present HypeNET, an LSTM-based method for hypernymy detection. We first focus on improving path representation (Section 3.1), and then integrate distributional signals into our network, resulting in a combined method (Section 3.2).

3.1 Path-based Network

Similarly to prior work, we represent each dependency path as a sequence of edges that leads from x to y in the dependency tree.² Each edge contains the lemma and part-of-speech tag of the source node, the dependency label, and the edge direction between two subsequent nodes. We denote each edge as *lemma/POS/dep/dir*. See figure 1 for an illustration.

Rather than treating an entire dependency path as a single feature, we encode the sequence of edges using a long short-term memory (LSTM) network. The vectors obtained for the different paths of a given (x, y) pair are pooled, and the resulting vector is used for classification. Figure 2 depicts the overall network structure, which is described below.

Edge Representation We represent each edge by the concatenation of its components’ vectors:

$$\vec{v}_e = [\vec{v}_l, v_{pos}, v_{dep}, v_{dir}]$$

where $\vec{v}_l, v_{pos}, v_{dep}, v_{dir}$ represent the embedding vectors of the lemma, part-of-speech, dependency label and dependency direction (along the path from x to y), respectively.

Path Representation For a path p composed of edges e_1, \dots, e_k , the edge vectors $\vec{v}_{e_1}, \dots, \vec{v}_{e_k}$ are fed in order to an LSTM encoder, resulting in a vector \vec{o}_p representing the entire path p . The LSTM architecture is effective at capturing temporal patterns in sequences. We expect the training procedure to drive the LSTM encoder to focus on parts of the path that are more informative for the classification task while ignoring others.

²Like Snow et al. (2004), we added for each path, additional paths containing single daughters of x or y not already contained in the path, referred by Snow et al. (2004) as “satellite edges”. This enables including paths like *Such Y as X*, in which the word “such” is not in the path between x and y .

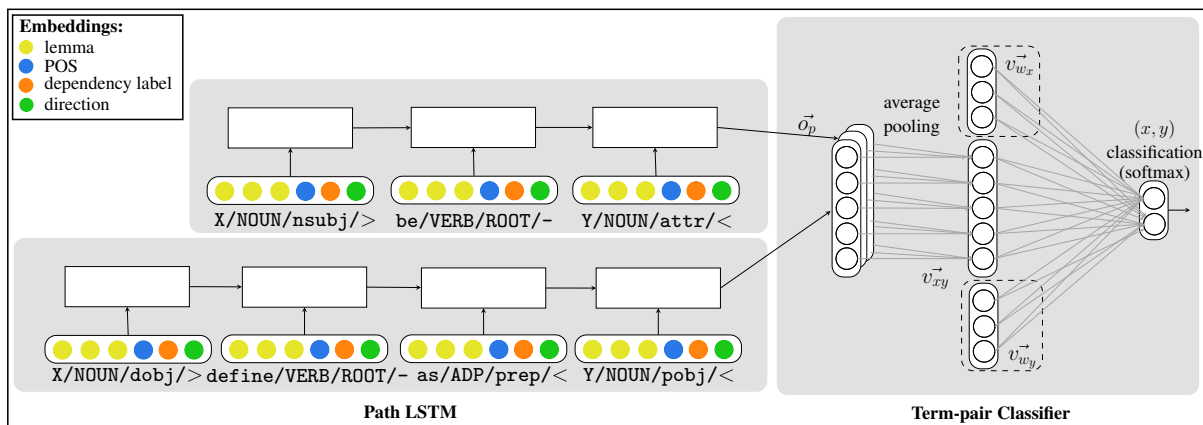


Figure 2: An illustration of term-pair classification. Each term-pair is represented by several paths. Each path is a sequence of edges, and each edge consists of four components: lemma, POS, dependency label and dependency direction. Each edge vector is fed in sequence into the LSTM, resulting in a path embedding vector \vec{o}_p . The averaged path vector becomes the term-pair’s feature vector, used for classification. The dashed v_{w_x} , v_{w_y} vectors refer to the integrated network described in Section 3.2.

Term-Pair Classification Each (x, y) term-pair is represented by the multiset of lexico-syntactic paths that connected x and y in the corpus, denoted as $paths(x, y)$, while the supervision is given for the term pairs. We represent each (x, y) term-pair as the weighted-average of its path vectors, by applying average pooling on its path vectors, as follows:

$$\vec{v}_{xy} = \vec{v}_{paths(x,y)} = \frac{\sum_{p \in paths(x,y)} f_{p,(x,y)} \vec{o}_p}{\sum_{p \in paths(x,y)} f_{p,(x,y)}} \quad (1)$$

where $f_{p,(x,y)}$ is the frequency of p in $paths(x, y)$. We then feed this path vector to a single-layer network that performs binary classification to decide whether y is a hypernym of x .

$$c = softmax(W \cdot v_{xy}) \quad (2)$$

c is a 2-dimensional vector whose components sum to 1, and we classify a pair as positive if $c[1] > 0.5$.

Implementation Details To train the network, we used PyCNN.³ We minimize the cross entropy loss using gradient-based optimization, with mini-batches of size 10 and the Adam update rule (Kingma and Ba, 2014). Regularization is applied by a dropout on each of the components’ embeddings. We tuned the hyper-parameters (learning rate and dropout rate) on the validation set (see the appendix for the hyper-parameters values).

We initialized the lemma embeddings with the pre-trained GloVe word embeddings (Pennington et al., 2014), trained on Wikipedia. We tried both

the 50-dimensional and 100-dimensional embedding vectors and selected the ones that yield better performance on the validation set.⁴ The other embeddings, as well as out-of-vocabulary lemmas, are initialized randomly. We update all embedding vectors during training.

3.2 Integrated Network

The network presented in Section 3.1 classifies each (x, y) term-pair based on the paths that connect x and y in the corpus. Our goal was to improve upon previous path-based methods for hypernymy detection, and we show in Section 6 that our network indeed outperforms them. Yet, as path-based and distributional methods are considered complementary, we present a simple way to integrate distributional features in the network, yielding improved performance.

We extended the network to take into account distributional information on each term. Inspired by the supervised distributional concatenation method (Baroni et al., 2012), we simply concatenate x and y word embeddings to the (x, y) feature vector, redefining v_{xy} :

$$v_{xy} = [v_{w_x}, \vec{v}_{paths(x,y)}, v_{w_y}] \quad (3)$$

where v_{w_x} and v_{w_y} are x and y ’s word embeddings, respectively, and $\vec{v}_{paths(x,y)}$ is the averaged path vector defined in equation 1. This way, each (x, y) pair is represented using both the distributional features of x and y , and their path-based features.

⁴Higher-dimensional embeddings seem not to improve performance, while hurting the training runtime.

³<https://github.com/clab/cnn>

resource	relations
WordNet	instance hypernym, hypernym
DBPedia	type
Wikidata	subclass of, instance of
Yago	subclass of

Table 1: Hypernymy relations in each resource.

4 Dataset

4.1 Creating Instances

Neural networks typically require a large amount of training data, whereas the existing hypernymy datasets, like BLESS (Baroni and Lenci, 2011), are relatively small. Therefore, we followed the common methodology of creating a dataset using distant supervision from knowledge resources (Snow et al., 2004; Riedel et al., 2013). Following Snow et al. (2004), who constructed their dataset based on WordNet hypernymy, and aiming to create a larger dataset, we extract hypernymy relations from several resources: WordNet (Fellbaum, 1998), DBPedia (Auer et al., 2007), Wikidata (Vrandečić, 2012) and Yago (Suchanek et al., 2007).

All instances in our dataset, both positive and negative, are pairs of terms that are directly related in at least one of the resources. These resources contain thousands of relations, some of which indicate hypernymy with varying degrees of certainty. To avoid including questionable relation types, we consider as denoting positive examples only indisputable hypernymy relations (Table 1), which we manually selected from the set of hypernymy indicating relations in Shwartz et al. (2015).

Term-pairs related by other relations (including hyponymy), are considered as negative instances. Using related rather than random term-pairs as negative instances tests our method’s ability to distinguish between hypernymy and other kinds of semantic relatedness. We maintain a ratio of 1:4 positive to negative pairs in the dataset.

Like Snow et al. (2004), we include only term-pairs that have joint occurrences in the corpus, requiring at least two different dependency paths for each pair.

4.2 Random and Lexical Dataset Splits

As our primary dataset, we perform standard random splitting, with 70% train, 25% test and 5% validation sets.

As pointed out by Levy et al. (2015), supervised distributional lexical inference methods tend

	train	test	val	all
random split	49,475	17,670	3,534	70,679
lexical split	20,335	6,610	1,350	28,295

Table 2: The number of instances in each dataset.

to perform “lexical memorization”, i.e., instead of learning a relation between the two terms, they mostly learn an independent property of a single term in the pair: whether it is a “prototypical hypernym” or not. For instance, if the training set contains term-pairs such as (*dog*, *animal*), (*cat*, *animal*), and (*cow*, *animal*), all annotated as positive examples, the algorithm may learn that *animal* is a prototypical hypernym, classifying any new (*x*, *animal*) pair as positive, regardless of the relation between *x* and *animal*. Levy et al. (2015) suggested to split the train and test sets such that each will contain a distinct vocabulary (“lexical split”), in order to prevent the model from overfitting by lexical memorization.

To investigate such behaviors, we present results also for a lexical split of our dataset. In this case, we split the train, test and validation sets such that each contains a distinct vocabulary. We note that this differs from Levy et al. (2015), who split only the train and the test sets, and dedicated a subset of the train for validation. We chose to deviate from Levy et al. (2015) because we noticed that when the validation set contains terms from the train set, the model is rewarded for lexical memorization when tuning the hyper-parameters, consequently yielding suboptimal performance on the lexically-distinct test set. When each set has a distinct vocabulary, the hyper-parameters are tuned to avoid lexical memorization and are likely to perform better on the test set. We tried to keep roughly the same 70/25/5 ratio in our lexical split.⁵ The sizes of the two datasets are shown in Table 2.

Indeed, training a model on a lexically split dataset may result in a more general model, that can better handle pairs consisting of two unseen terms during inference. However, we argue that in the common applied scenario, the inference involves an unseen pair (*x*, *y*), in which *x* and/or *y* have already been observed separately. Models trained on a random split may introduce the model with a term’s “prior probability” of being a hypernym or a hyponym, and this information can be exploited beneficially at inference time.

⁵The lexical split discards many pairs consisting of cross-set terms.

path
X/NOUN/dobj/> establish/VERB/ROOT/- as/ADP/prep/< Y/NOUN/pobj/<
X/NOUN/dobj/> VERB as/ADP/prep/< Y/NOUN/pobj/<
X/NOUN/dobj/> * as/ADP/prep/< Y/NOUN/pobj/<
X/NOUN/dobj/> establish/VERB/ROOT/- ADP Y/NOUN/pobj/<
X/NOUN/dobj/> establish/VERB/ROOT/- * Y/NOUN/pobj/<

Table 3: Example generalizations of *X was established as Y*.

5 Baselines

We compare HypeNET with several state-of-the-art methods for hypernymy detection, as described in Section 2: path-based methods (Section 5.1), and distributional methods (Section 5.2). Due to different works using different datasets and corpora, we replicated the baselines rather than comparing to the reported results.

We use the Wikipedia dump from May 2015 as the underlying corpus of all the methods, and parse it using spaCy.⁶ We perform model selection on the validation set to tune the hyper-parameters of each method.⁷ The best hyper-parameters are reported in the appendix.

5.1 Path-based Methods

Snow We follow the original paper, and extract all shortest paths of four edges or less between terms in a dependency tree. Like Snow et al. (2004), we add paths with “satellite edges”, i.e., single words not already contained in the dependency path, which are connected to either X or Y, allowing paths like *such Y as X*. The number of distinct paths was 324,578. We apply χ^2 feature selection to keep only the 100,000 most informative paths and train a logistic regression classifier.

Generalization We also compare our method to a baseline that uses generalized dependency paths. Following PATTY’s approach to generalizing paths (Nakashole et al., 2012), we replace edges with their part-of-speech tags as well as with wild cards. We generate the powerset of all possible generalizations, including the original paths. See Table 3 for examples. The number of features after generalization went up to 2,093,220. Similarly to the first baseline, we apply feature selection, this time keeping the 1,000,000 most informative paths, and train a logistic regression classifier over the generalized paths.⁸

⁶<https://spacy.io/>

⁷We applied grid search for a range of values, and picked the ones that yield the highest F_1 score on the validation set.

⁸We also tried keeping the 100,000 most informative paths, but the performance was worse.

5.2 Distributional Methods

Unsupervised SLQS (Santus et al., 2014) is an entropy-based measure for hypernymy detection, reported to outperform previous state-of-the-art unsupervised methods (Weeds and Weir, 2003; Kotlerman et al., 2010). The original paper was evaluated on the BLESS dataset (Baroni and Lenci, 2011), which consists of mostly frequent words. Applying the vanilla settings of SLQS on our dataset, that contains also rare terms, resulted in low performance. Therefore, we received assistance from Enrico Santus, who kindly provided the results of SLQS on our dataset after tuning the system as follows.

The validation set was used to tune the threshold for classifying a pair as positive, as well as the maximum number of each term’s most associated contexts (N). In contrast to the original paper, in which the number of each term’s contexts is fixed to N , in this adaptation it was set to the minimum between the number of contexts with LMI score above zero and N . In addition, the SLQS scores were not multiplied by the cosine similarity scores between terms, and terms were lemmatized prior to computing the SLQS scores, significantly improving recall.

As our results suggest, while this method is state-of-the-art for unsupervised hypernymy detection, it is basically designed for classifying specificity level of related terms, rather than hypernymy in particular.

Supervised To represent term-pairs with distributional features, we tried several state-of-the-art methods: concatenation $\vec{x} \oplus \vec{y}$ (Baroni et al., 2012), difference $\vec{y} - \vec{x}$ (Roller et al., 2014; Weeds et al., 2014), and dot-product $\vec{x} \cdot \vec{y}$. We downloaded several pre-trained embeddings (Mikolov et al., 2013; Pennington et al., 2014) of different sizes, and trained a number of classifiers: logistic regression, SVM, and SVM with RBF kernel, which was reported by Levy et al. (2015) to perform best in this setting. We perform model selection on the validation set to select the best vectors, method and regularization factor (see the appendix).

method		random split			lexical split		
		precision	recall	F_1	precision	recall	F_1
Path-based	Snow	0.843	0.452	0.589	0.760	0.438	0.556
	Snow + Gen	0.852	0.561	0.676	0.759	0.530	0.624
	HypeNET Path-based	0.811	0.716	0.761	0.691	0.632	0.660
Distributional	SLQS (Santus et al., 2014)	0.491	0.737	0.589	0.375	0.610	0.464
	Best supervised (concatenation)	0.901	0.637	0.746	0.754	0.551	0.637
Combined	HypeNET Integrated	0.913	0.890	0.901	0.809	0.617	0.700

Table 4: Performance scores of our method compared to the path-based baselines and the state-of-the-art distributional methods for hypernymy detection, on both variations of the dataset – with lexical and random split to train / test / validation.

6 Results

Table 4 displays performance scores of HypeNET and the baselines. *HypeNET Path-based* is our path-based recurrent neural network model (Section 3.1) and *HypeNET Integrated* is our combined method (Section 3.2). Comparing the path-based methods shows that generalizing paths improves recall while maintaining similar levels of precision, reassessing the behavior found in Nakashole et al. (2012). *HypeNET Path-based* outperforms both path-based baselines by a significant improvement in recall and with slightly lower precision. The recall boost is due to better path generalization, as demonstrated in Section 7.1.

Regarding distributional methods, the unsupervised SLQS baseline performed slightly worse on our dataset. The low precision stems from its inability to distinguish between hypernyms and meronyms, which are common in our dataset, causing many false positive pairs such as (*zabrze, poland*) and (*kibbutz, israel*). We sampled 50 false positive pairs of each dataset split, and found that 38% of the false positive pairs in the random split and 48% of those in the lexical split were holonym-meronym pairs.

In accordance with previously reported results, the supervised embedding-based method is the best performing baseline on our dataset as well. *HypeNET Path-based* performs slightly better, achieving state-of-the-art results. Adding distributional features to our method shows that these two approaches are indeed complementary. On both dataset splits, the performance differences between *HypeNET Integrated* and *HypeNET Path-based*, as well as the supervised distributional method, are substantial, and statistically significant with p-value of 1% (paired t-test).

We also reassess that indeed supervised distributional methods perform worse on a lexical split (Levy et al., 2015). We further observe a similar reduction when using HypeNET, which is not a re-

sult of lexical memorization, but rather stems from over-generalization (Section 7.1).

7 Analysis

7.1 Qualitative Analysis of Learned Paths

We analyze HypeNET’s ability to generalize over path structures, by comparing prominent indicative paths which were learned by each of the path-based methods. We do so by finding high-scoring paths that contributed to the classification of true-positive pairs in the dataset. In the path-based baselines, these are the highest-weighted features as learned by the logistic regression classifier. In the LSTM-based method, it is less straightforward to identify the most indicative paths. We assess the contribution of a certain path p to classification by regarding it as the only path that appeared for the term-pair, and compute its TRUE label score from the class distribution: $\text{softmax}(W \cdot v_{xy})[1]$, setting $v_{xy} = [\vec{0}, \vec{\sigma}_p, \vec{0}]$.

A notable pattern is that Snow’s method learns specific paths, like *X is Y from* (e.g. *Megadeth is an American thrash metal band from Los Angeles*). While Snow’s method can only rely on verbatim paths, limiting its recall, the generalized version of Snow often makes coarse generalizations, such as *X VERB Y from*. Clearly, such a path is too general, and almost any verb assigned to it results in a non-indicative path (e.g. *X take Y from*). Efforts by the learning method to avoid such generalization, again, lower the recall. HypeNET provides a better midpoint, making fine-grained generalizations by learning additional semantically similar paths such as *X become Y from* and *X remain Y from*. See table 5 for additional example paths which illustrate these behaviors.

We also noticed that while on the random split our model learns a range of specific paths such as *X is Y published* (learned for e.g. $Y=\text{magazine}$) and *X is Y produced* ($Y=\text{film}$), in the lexical split it only learns the general *X is Y* path for these re-

method	path	example text
Snow	X/NOUN/nsubj/> be/VERB/ROOT/- Y/NOUN/attr/< direct/VERB/acl/>	<i>Eyeball</i> is a 1975 Italian-Spanish <i>film</i> directed by Umberto Lenzi
	X/NOUN/nsubj/> be/VERB/ROOT/- Y/NOUN/attr/< publish/VERB/acl/>	<i>Allure</i> is a U.S. women’s beauty <i>magazine</i> published monthly
Snow + Gen	X/NOUN/compound/> NOUN* be/VERB/ROOT/- Y/NOUN/attr/< base/VERB/acl/>	<i>Calico</i> Light Weapons Inc. (CLWS) is an American privately held manufacturing <i>company</i> based in Cornelius, Oregon
	X/NOUN/compound/> NOUN Y/NOUN/compound/<	<i>Weston Town</i> Council
HypeNET Integrated	X/NOUN/nsubj/> be/VERB/ROOT/- Y/NOUN/attr/< (release direct produce write)/VERB/acl/>	<i>Blinky</i> is a 1923 American comedy <i>film</i> directed by Edward Sedgwick
	X/NOUN/compound/> (association co. company corporation foundation group inc. international limited ltd.)/NOUN/nsubj/> be/VERB/ROOT/- Y/NOUN/attr/< ((create found headquarter own specialize)/VERB/acl/>)?	<i>Retalix</i> Ltd. is a software <i>company</i>

Table 5: Examples of indicative paths learned by each method, with corresponding true positive term-pairs from the random split test set. Hypernyms are marked red and hyponyms are marked blue.

Relation	%
synonymy	21.37%
hyponymy	29.45%
holonymy / meronymy	9.36%
hypernymy-like relations	21.03%
other relations	18.77%

Table 6: Distribution of relations holding between each pair of terms in the resources among false positive pairs.

lations. We note that *X is Y* is a rather “noisy” path, which may occur in ad-hoc contexts without indicating generic hypernymy relations (e.g. *chocolate is a big problem* in the context of children’s health). While such a model may identify hypernymy relations between unseen terms, based on general paths, it is prone to over-generalization, hurting its performance, as seen in Table 4. As discussed in § 4.2, we suspect that this scenario, in which both terms are unseen, is usually not common enough to justify this limiting training setup.

7.2 Error Analysis

False Positives We categorized the false positive pairs on the random split according to the relation holding between each pair of terms in the resources used to construct the dataset. We grouped several semantic relations from different resources to broad categories, e.g. *synonym* includes also *alias* and *Wikipedia redirection*. Table 6 displays the distribution of semantic relations among false positive pairs.

More than 20% of the errors stem from confusing synonymy with hypernymy, which are known to be difficult to distinguish.

An additional 30% of the term-pairs are re-

versed hypernym-hyponym pairs (*y* is a hyponym of *x*). Examining a sample of these pairs suggests that they are usually near-synonyms, i.e., it is not that clear whether one term is truly more general than the other or not. For instance, *fiction* is annotated in WordNet as a hypernym of *story*, while our method classified *fiction* as its hyponym.

A possible future research direction might be to quite simply extend our network to classify term-pairs simultaneously to multiple semantic relations, as in Pavlick et al. (2015). Such a multi-class model can hopefully better distinguish between these similar semantic relations.

Another notable category is hypernymy-like relations: these are other relations in the resources that could also be considered as hypernymy, but were annotated as negative due to our restrictive selection of only indisputable hypernymy relations from the resources (see Section 4.1). These include instances like (*Goethe, occupation, novelist*) and (*Homo, subdivisionRanks, species*).

Lastly, other errors made by the model often correspond to term-pairs that co-occur very few times in the corpus, e.g. *xebec*, a studio producing Anime, was falsely classified as a hyponym of *anime*.

False Negatives We sampled 50 term-pairs that were falsely annotated as negative, and analyzed the major (overlapping) types of errors (Table 7).

Most of these pairs had only few co-occurrences in the corpus. This is often either due to infrequent terms (e.g. *cbc.ca*), or a rare sense of *x* in which the hypernymy relation holds (e.g. (*night*,

	Error Type	%
1	low statistics	80%
2	infrequent term	36%
3	rare hyponym sense	16%
4	annotation error	8%

Table 7: (Overlapping) categories of false negative pairs: (1) x and y co-occurred less than 25 times (average co-occurrences for true positive pairs is 99.7). (2) Either x or y is infrequent. (3) The hypernymy relation holds for a rare sense of x . (4) (x, y) was incorrectly annotated as positive.

play) holding for “Night”, a dramatic sketch by Harold Pinter). Such a term-pair may have too few hypernymy-indicating paths, leading to classifying it as negative.

8 Conclusion

We presented HypeNET: a neural-networks-based method for hypernymy detection. First, we focused on improving path representation using LSTM, resulting in a path-based model that performs significantly better than prior path-based methods, and matches the performance of the previously superior distributional methods. In particular, we demonstrated that the increase in recall is a result of generalizing semantically-similar paths, in contrast to prior methods, which either make no generalizations or over-generalize paths.

We then extended our network by integrating distributional signals, yielding an improvement of additional 14 F_1 points, and demonstrating that the path-based and the distributional approaches are indeed complementary.

Finally, our architecture seems straightforwardly applicable for multi-class classification, which, in future work, could be used to classify term-pairs to multiple semantic relations.

Acknowledgments

We would like to thank Omer Levy for his involvement and assistance in the early stage of this project and Enrico Santus for helping us by computing the results of SLQS (Santus et al., 2014) on our dataset.

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).

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.
- Marco Baroni and Alessandro Lenci. 2011. How we blessed distributional semantic evaluation. In *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 *EACL*, pages 23–32.
- Claudia Borg, Mike Rosner, and Gordon Pace. 2009. Evolutionary algorithms for definition extraction. In *Proceedings of the 1st Workshop on Definition Extraction*, pages 26–32.
- Andrew Carlson, Justin Betteridge, Bryan Kisiel, Burr Settles, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Toward an architecture for never-ending language learning. In *AAAI*.
- Christiane Fellbaum. 1998. *WordNet*. Wiley Online Library.
- Katrin Fundel, Robert Küffner, and Ralf Zimmer. 2007. Relexrelation extraction using dependency parse trees. *Bioinformatics*, pages 365–371.
- Marti A Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *ACL*, pages 539–545.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2009. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *SemEval*, pages 94–99.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, pages 1735–1780.
- Nobuhiro Kaji and Masaru Kitsuregawa. 2008. Using hidden markov random fields to combine distributional and pattern-based word clustering. In *COLING*, pages 401–408.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2010. Directional distributional similarity for lexical inference. *NLE*, pages 359–389.
- Zornitsa Kozareva and Eduard Hovy. 2010. A semi-supervised method to learn and construct taxonomies using the web. In *EMNLP*, pages 1110–1118.

- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015. Do supervised distributional methods really learn lexical inference relations. *NAACL*.
- Decang Lin. 1998. An information-theoretic definition of similarity. In *ICML*, pages 296–304.
- Yang Liu, Furu Wei, Sujian Li, Heng Ji, Ming Zhou, and Houfeng Wang. 2015. A dependency-based neural network for relation classification. *arXiv preprint arXiv:1507.04646*.
- 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.
- Shachar Mirkin, Ido Dagan, and Maayan Geffet. 2006. Integrating pattern-based and distributional similarity methods for lexical entailment acquisition. In *COLING and ACL*, pages 579–586.
- Ndapandula Nakashole, Gerhard Weikum, and Fabian Suchanek. 2012. Patty: a taxonomy of relational patterns with semantic types. In *EMNLP and CoNLL*, pages 1135–1145.
- Roberto Navigli and Paola Velardi. 2010. Learning word-class lattices for definition and hypernym extraction. In *ACL*, pages 1318–1327.
- Thien Huu Nguyen and Ralph Grishman. 2015. Combining neural networks and log-linear models to improve relation extraction. *arXiv preprint arXiv:1511.05926*.
- Ellie Pavlick, Johan Bos, Malvina Nissim, Charley Beller, Benjamin Van Durme, and Chris Callison-Burch. 2015. Adding semantics to data-driven paraphrasing. In *ACL*.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M Marlin. 2013. Relation extraction with matrix factorization and universal schemas. In *NAACL*.
- Laura Rimell. 2014. Distributional lexical entailment by topic coherence. In *EACL*, pages 511–519.
- Stephen Roller, Katrin Erk, and Gemma Boleda. 2014. Inclusive yet selective: Supervised distributional hypernymy detection. In *COLING*, pages 1025–1036.
- Enrico Santus, Alessandro Lenci, Qin Lu, and Sabine Schulte Im Walde. 2014. Chasing hypernyms in vector spaces with entropy. In *EACL*, pages 38–42.
- Vered Shwartz, Omer Levy, Ido Dagan, and Jacob Goldberger. 2015. Learning to exploit structured resources for lexical inference. In *CoNLL*, page 175.
- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2004. Learning syntactic patterns for automatic hypernym discovery. In *NIPS*.
- Rion Snow, Daniel Jurafsky, and Andrew Y Ng. 2006. Semantic taxonomy induction from heterogenous evidence. In *ACL*, pages 801–808.
- Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. Yago: a core of semantic knowledge. In *WWW*, pages 697–706.
- Peter D Turney. 2006. Similarity of semantic relations. *CL*, pages 379–416.
- Denny Vrandečić. 2012. Wikidata: A new platform for collaborative data collection. In *WWW*, pages 1063–1064.
- Julie Weeds and David Weir. 2003. A general framework for distributional similarity. In *EMLP*, pages 81–88.
- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. Learning to distinguish hypernyms and co-hyponyms. In *COLING*, pages 2249–2259.
- Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. Classifying relations via long short term memory networks along shortest dependency paths. In *EMNLP*.
- Yan Xu, Ran Jia, Lili Mou, Ge Li, Yunchuan Chen, Yangyang Lu, and Zhi Jin. 2016. Improved relation classification by deep recurrent neural networks with data augmentation. *arXiv preprint arXiv:1601.03651*.

Appendix A Best Hyper-parameters

Table 8 displays the chosen hyper-parameters of each method, yielding the highest F_1 score on the validation set.

	method	values
random split	Snow	regularization: L_2
	Snow + Gen	regularization: L_1
	LSTM	embeddings: GloVe-100-Wiki learning rate: $\alpha = 0.001$ dropout: $d = 0.5$
	SLQS	$N=100$, threshold = 0.000464
	Best Supervised	method: concatenation, classifier: SVM embeddings: GloVe-300-Wiki embeddings: GloVe-50-Wiki
	LSTM-Integrated	learning rate: $\alpha = 0.001$ word dropout: $d = 0.3$
lexical split	Snow	regularization: L_2
	Snow + Gen	regularization: L_2
	LSTM	embeddings: GloVe-50-Wiki learning rate: $\alpha = 0.001$ dropout: $d = 0.5$
	SLQS	$N=100$, threshold = 0.007629
	Best Supervised	method: concatenation, classifier: SVM embeddings: GloVe-100-Wikipedia embeddings: GloVe-50-Wiki
	LSTM-Integrated	learning rate: $\alpha = 0.001$ word dropout: $d = 0.3$

Table 8: The best hyper-parameters in every model.